Submitted to the Annals of Statistics

arXiv:1702.07283v1 [stat.ME] 23 Feb 2017

NON-PENALIZED VARIABLE SELECTION IN HIGH-DIMENSIONAL LINEAR MODEL SETTINGS VIA GENERALIZED FIDUCIAL INFERENCE

By Jonathan P Williams and Jan Hannig

University of North Carolina at Chapel Hill

Standard penalized methods of variable selection and parameter estimation rely on the magnitude of coefficient estimates to decide which variables to include in the final model. However, coefficient estimates are unreliable when the design matrix is collinear. To overcome this challenge an entirely new method of variable selection is presented within a generalized fiducial inference framework. This new procedure is able to effectively account for linear dependencies among subsets of covariates in a high-dimensional setting where p can grow almost exponentially in n, as well as in the classical setting where $p \leq n$.

It is shown that the procedure very naturally assigns small probabilities to subsets of covariates which include redundancies by way of explicit L_0 minimization. Furthermore, with a typical sparsity assumption, it is shown that the proposed method is consistent in the sense that the probability of the true sparse subset of covariates converges in probability to 1 as $n \to \infty$, or as $n \to \infty$ and $p \to \infty$. Very reasonable conditions are needed, and little restriction is placed on the class of 2^p possible subsets of covariates to achieve this consistency result.

1. Introduction. The high-dimensional regression setting poses a very difficult variable selection problem which becomes considerably more challenging in the presence of high collinearity among the explanatory variables. Consider the linear model

$$Y = X\beta + \sigma\vec{e} \text{ where } \vec{e} \sim \mathrm{N}_n(0, I_n),$$

where X is a $n \times p$ matrix, β is a $p \times 1$ vector (often believed to be sparse), and $\sigma > 0$. Assume that the columns of X have been scaled to have unit norm.

Standard penalized methods of variable selection and parameter estimation such as LASSO of Tibshirani (1996), SCAD of Fan and Li (2001), and

MSC 2010 subject classifications: 62J05,62F12,62A01

Keywords and phrases: variable selection, best subset selection, high dimensional regression, L_0 minimization, generalized fiducial inference

the Dantzig Selector of Candes and Tao (2007) rely on the magnitude of the coefficient estimates to decide which variables to include in the final model. However, regression coefficients lose identifiability when the observed covariates are linearly dependent because a theoretically true subset of observed covariates could be re-expressed as various linear combinations of the same covariates. Furthermore, in collinear design settings a true subset of covariates is not necessarily uniquely defined. This is problematic in applications where it is meaningful to be able to identify the true effects of the explanatory variables on the response variable.

To overcome such challenges an entirely new method of variable selection is presented in this paper within a *generalized fiducial inference* (GFI) framework (Hannig et al., 2016). The proposed method is developed on the idea that a true sparse subset of covariates for predicting a response should be defined with the property that there exists no smaller subset of covariates which can perform *predictively* within some neighborhood of the true model. This property will not hold if the true model is sufficiently correlated with other covariates being considered.

The variable selection procedure proposed in this paper is able to effectively account for linear dependencies among subsets of covariates in a high-dimensional setting where p can grow almost exponentially in n, as well as in the classical setting where $p \leq n$. It is shown that the procedure very naturally assigns small probabilities to subsets of covariates which include redundancies by way of explicit L_0 minimization. As a result, a probability distribution is defined whose mass is concentrated over all *least complex* subsets of covariates. The technique for the L_0 minimization is due to Bertsimas, King and Mazumder (2016).

Most notably is that the proposed variable selection procedure accomplishes the task of reducing model complexity in the sparse setting without relying on the magnitudes of the coefficient estimates; it is in no way akin to a penalized method of variable selection and parameter estimation. Furthermore, it is shown that the proposed method is consistent in the sense that the probability of the true sparse subset of covariates converges in probability to 1 as $n \to \infty$, or as $n \to \infty$ and $p \to \infty$. Very reasonable conditions are needed, and little restriction is placed on the class of 2^p possible subsets of covariates to achieve this consistency result.

Inspiration for the variable selection procedure proposed in this paper came from Johnson and Rossell (2012). They introduce the idea of using nonlocal prior densities on the model parameters in high-dimensional regression settings, and show that Bayesian approaches with such priors will assign posterior probability 1 to the true model, asymptotically. They also show

that Bayesian approaches to variable selection based on local prior specifications lack this consistency result because they assign too much mass around zero to redundant parameter estimates. The idea behind nonlocal priors fits naturally into the GFI setting. Indeed, GFI provides natural tools for developing a method which assigns negligible probabilities to candidate subsets of covariates which include redundancies.

Other relevant work on variable selection in the Bayesian framework includes the credible set approach of Bondell and Reich (2012), as well as Narisetty and He (2014) which propose a method based on shrinking and diffusing coefficient priors in which the variance of the priors are sample size dependent. Ghosh and Ghattas (2015) provide important insights into complications in Bayesian variable selection when the design matrix exhibits strong collinearity. Namely, the size of the sample space (2^p) is often too large to compute all model probabilities, and even typically larger than can reasonably be sampled by MCMC methods. They note problems with selecting candidate models by choosing the highest probability model or the median probability model. For instance, high-dimensional settings often lead to arbitrarily small probabilities for all models simply because there are too many models to consider. Ghosh and Ghattas (2015) also discuss problems with relying on marginal inclusion probabilities, and rather argue that joint summaries of subset of covariates are more robust to collinearity. It is in this spirit that the proposed GFI variable selection methodology is introduced in this paper.

The paper is organized as follows. Section 2 serves to introduce the general methodology and computational algorithm. The conditions needed for the main results, and the main results are presented and discussed in Section 3. Proofs are organized in the appendix. Simulation results are given in Section 4, and an application of the proposed method to a high-dimensional, real data set involving gene expressions is explored in Section 5.

2. Methodology. The GFI framework expresses observed data $Y = G((\beta, \sigma), \vec{e})$ for some function G known as the data generating function, with (β, σ) being the parameters and $\vec{e}$ being a random vector with completely known distribution independent of any parameters. In the linear model setting $G((\beta, \sigma), \vec{e}) = X\beta + \sigma\vec{e}$. The next step is to apply the 'switching principle' which defines a distribution on the parameters by inverting the data generating equation for an observed data vector y. That is, $(\beta^\star, \sigma^\star) = G^{-1}(y, \vec{e}^\star)$ where the distribution of $\vec{e}^\star$ is the distribution of $\vec{e}$ conditional on the set on which the inverse exists, and $\beta^\star$ and $\sigma^\star$ are used to distinguish the parameters as random, instead of fixed. The details can be found in a recent

review article by Hannig et al. (2016).

The nonlocal prior approach of Johnson and Rossell (2012) represents a new foundation for thinking about variable selection problems. Many variable selection procedures attempt to facilitate the computation of probabilities of all 2^p possible subsets of covariates, and in falling short resort to explicitly restricting the class of subsets to consider in order to achieve a reasonable consistency result. However, many of the candidate subsets of covariates are redundant in the sense that they contain unnecessary covariates which contribute little to no additional explanatory power. The implementation of a nonlocal prior from Johnson and Rossell (2012) defines a conditional posterior probability distribution on the class of subsets of covariates $M \subset \{1, \ldots, p\}$ such that M contains only covariates with nonzero coefficients. This yields a probability distribution with mass concentrated only over subsets M without redundancies, but results in an unsettling interpretation for the subsets with zero posterior probability. This is because the nonlocal priors give zero density to the true coefficient value when the true coefficient value is zero.

This motivates an entirely new approach to variable selection which fits naturally with the idea of a nonlocal prior. The fundamental difference is to treat subsets of covariates holistically rather than relying on coefficient estimates, and to re-define the notion of redundant covariates. Accordingly, define a *least complex* subsets of covariates to be one such that no *smaller* subset of covariates exists which performs *predictively within* ε of the subset of covariates being considered. Then, when applying the 'switching principle' on the data generating equation, we consider only *least complex* subsets. This results in a probability distribution with mass concentrated over subsets M without redundancies, and all other subsets will have negligibly small probabilities without imposing unnatural prior restrictions. Note that the definition of a *least complex* subset of covariates inherently addresses the collinearity within a subset of covariates. The discussion to follow is focused on making this definition precise.

For a given subset M of the p covariates being considered define their *prediction* to be $X_M\beta^\star_M$ where X_M is the sub-matrix of the design matrix X consisting only of those columns corresponding to the covariates in M, and $\beta^\star_M$ is the random vector of coefficients corresponding to the covariates in M. Define a *smaller* subset of the p possible covariates to be any subset with less than $|M|$ covariates, where $|M|$ denotes the cardinality of M. Accordingly, the difference in prediction between M and all smaller subsets is given by the minimum distance between $X_M\beta^\star_M$ and Xb over all $b \in \mathbb{R}^p$ with $\|b\|_0 \leq |M| - 1$.

The chosen measure of distance will be $\frac{1}{2}\|X^T(X_M\beta^\star_M - Xb)\|_2^2$ which has been adapted from Candes and Tao (2007). However, some clarification is needed at this point. The Dantzig Selector from Candes and Tao (2007) deals with the error $\|X^T(y - Xb)\|_\infty$ over $b \in \mathbb{R}^p$ where y is a random response vector. This is very different from using $X_M\beta^\star_M$ in place of y because the former results in a noiseless measure of distance. In other words, if one of the components of $\beta^\star_M$ is zero, then $X_M\beta^\star_M$ can be exactly recovered over the space of all $b \in \mathbb{R}^p$ with $\|b\|_0 \leq |M| - 1$.

There are a two main reasons to use $X^T(X_M\beta^\star_M - Xb)$ versus simply the difference in predictions, $(X_M\beta^\star_M - Xb)$. One reason is that $X^T(X_M\beta^\star_M - Xb)$ accounts for the difference in predictions as well as it's correlation with the explanatory data, as discussed in Berk (2008). If the difference in predictions is small but is highly correlated with the design matrix, then it is likely that the smaller subset of covariates is unable to account for the effect of one or more of the covariates in M. Thus, using $X^T(X_M\beta^\star_M - Xb)$ instead of just the difference in predictions is a method of controlling for potential omitted variable effects which could incorrectly find a close fitting subset of covariates that is *less complex* than M. Another advantage of $X^T(X_M\beta^\star_M - Xb)$ as a measure of distance is that it is invariant under orthogonal transformations of the design matrix, as pointed out in Candes and Tao (2007).

With the distance now chosen for measuring how close $X_M\beta^\star_M$ can be matched in prediction accuracy by a *less complex* subset of covariates, the decision function for defining what is considered a *close prediction* will be denoted by

$$h^\varepsilon_M(\beta^\star_M) := 1\Big\{\frac{1}{2}\|X^T(X_M\beta^\star_M - Xb_{min})\|_2^2 \geq \varepsilon_M(n,p)\Big\}, \tag{1}$$

where b_{min} solves

$$\min_{b\in\mathbb{R}^p} \frac{1}{2}\|X^T(X_M\beta^\star_M - Xb)\|_2^2 \quad \text{subject to} \quad \|b\|_0 \leq |M| - 1.$$

Note that if the covariates in M are linearly dependent, then $h^\varepsilon_M(\beta^\star_M) = 0$ for all $\beta^\star_M$. The subtlety in defining this decision rule is in choosing an appropriate ε which exceeds the defined difference in predictions when a given subset M is not *least complex*, and does not otherwise. Since, our chosen measure of distance naturally grows with $p_M := |M|$, the size of the model, ε must depend on the subset M under consideration. Additionally, the defined differences in predictions for a given M are more identifiable when more data observations are available, but predictions can be matched closer when there are more covariates to choose from. Therefore, ε must also

be a function of n and p, and so will be denoted by $\varepsilon_M(n,p)$. The choice of $\varepsilon_M(n,p)$ is a major focus of Section 3 where the main results of the paper are presented, and a good default choice can be found in (4).

The function h_M^ε manifests itself as a multiplicative factor in the generalized fiducial mass function for each M as a result of inverting the data generating equation, after first conditioning the distribution of $\vec{e}^\star$ on the inverse existing for $\beta_M^\star$ within the class of *least complex* subsets.

It is at this point that we incorporate the particulars of the generalized fiducial inference framework, but the reader should observe that an analogous distribution could have been constructed in a Bayesian framework. As in Hannig et al. (2016), the generalized fiducial distribution function necessitates a choice of the Jacobian term. A natural choice is to use the L_2 norm which gives

$$J(y,(\beta^\star,\sigma^\star)) = D\Big[\Big(X, \frac{y - X\beta^\star}{\sigma^\star}\Big)\Big] = \sigma^{\star -1} |\det(X^TX)|^{\frac{1}{2}} RSS^{\frac{1}{2}}.$$

Then conditioning on the class of all *least complex* subsets, the marginal generalized fiducial mass function for each M is

$$r(M|y) := \frac{\int h_M^\varepsilon(\beta_M^\star) f_M(y|\beta_M^\star,\sigma_M^\star) J_M(y,(\beta_M^\star,\sigma_M^\star))\ d(\sigma_M^\star,\beta_M^\star)}{\sum_{j=1}^{p}\sum_{|M|=j}\int h_M^\varepsilon(\beta_M^\star) f_M(y|\beta_M^\star,\sigma_M^\star) J_M(y,(\beta_M^\star,\sigma_M^\star))\ d(\sigma_M^\star,\beta_M^\star)}.$$

Here, and where it is notationally more convenient let $p_M := |M|$. Denoting the normalizing constant by C, the mass function simplifies to

$$r(M|y) = C \cdot \pi^{\frac{p_M}{2}} \Gamma\Big(\frac{n-p_M}{2}\Big) RSS_M^{-(\frac{n-p_M-1}{2})} \mathrm{E}[h_M^\varepsilon(\beta_M^\star)], \tag{2}$$

where the expectation is taken with respect to the location scale T distribution

$$t_{n-p_M}\Big((X_M^TX_M)^{-1}X_M^Ty, \frac{RSS_M}{n-p_M}(X_M^TX_M)^{-1}\Big). \tag{3}$$

As desired, $r(M|y)$ defines a probability distribution whose mass is concentrated over all *least complex* subset of covariates, M.

2.1. *Remarks on computation.* With the methodology now established, the obvious question is whether or not $r(M|y)$ can be efficiently computed. It can, using two recent computational tools. The first is an explicit L_0 minimization algorithm introduced in Bertsimas, King and Mazumder (2016) used to evaluate h_M^ε. The authors state that their algorithm borrows ideas

from projected gradient descent and methods in first-order convex optimization, and solves problems of the form

$$\min_{b\in\mathbb{R}^p} g(b) \quad \text{subject to} \quad \|b\|_0 \leq \kappa,$$

where $g(b) \geq 0$ is convex and has Lipschitz continuous gradient:

$$\|\nabla g(b) - \nabla g(\widetilde{b})\|_2 \leq l\|b - \widetilde{b}\|_2.$$

Within the broader context of applications of this algorithm, the practitioner may have little to no prior knowledge about how to specify κ. However, h^{ε}_M very naturally defines $\kappa := |M| - 1$ as in (1). Further, for the reader that is skeptical about an algorithm which claims to perform L_0 minimization, observe that in (1) it is not necessary to actually obtain $\min_{b\in\mathbb{R}^p} \frac{1}{2}\|X^T(X_M\beta^{\star}_M - Xb)\|_2^2$, but only to determine if it is larger or smaller than $\varepsilon_M(n,p)$. Thus, our algorithm will terminate once/if it finds an appropriate b such that $h^{\varepsilon}_M(\beta^{\star}_M) = 0$, rather than having to find the argument minimum. See Appendix A.1 and Algorithm A.1 for the details.

The second tool for computation is the Grouped Independence Metropolis Hastings (GIMH) algorithm from Andrieu and Roberts (2009), but originally introduced in Beaumont (2003). There is no obvious closed form expression for the probability mass function (2), and in fact, the mass function cannot even be expressed up to a normalizing constant because of the expectation in the expression. So standard MCMC techniques do not apply. However, Andrieu and Roberts (2009) describe GIMH as appropriate for situations in which the desired target distribution is analytically intractable or too complex to evaluate, but introducing a latent variable leads to analytical expressions or easier implementation. This is precisely the case with $r(M|y)$ which can be evaluated with Algorithm A.1 if the expectation is omitted, introducing a latent location scale T vector. See Appendix A.2 and Algorithm A.2 for the details.

3. Theoretical results. The main objective of this section is to show that the generalized fiducial probability of the true model will converge to 1 as the sample size, n, and number of covariates, p, are taken to infinity. First, the necessary assumptions and lemmas are stated. In the notation that follows denote the true model by M_T, define $p_T := |M_T|$, denote the true coefficient vector by $\beta^{\star}_{M_T}$, and define $\mu_T := X_{M_T}\mathrm{E}(\widehat{\beta}_{M_T})$. The quantity $\mathrm{E}(\widehat{\beta}_M)$ represents the expectation taken over the data, y, of $\widehat{\beta}_M := (X_M^T X_M)^{-1}X_M^T y$. Arguably, a '$y$' subscript should be placed on this expectation, but it is suppressed for notational conciseness. M_T is defined

as the *least complex* sparse subset of covariates from which the data y are generated. Recall that a *least complex* subset M is a subset such that no *smaller* subset of covariates is *predictively* within some neighborhood of M, that is $h_M^\varepsilon(\beta_M^\star) = 1$. Recall the definition of $h_M^\varepsilon(\beta_M^\star)$ from (1).

While h_M^ε cannot be evaluated analytically for many subsets M, it is zero for all M with $|M| > n$. If M contains more than n covariates, then the covariates (vectors of length n) in M are necessarily linearly dependent. In that case, there certainly exists a *less complex* subset of covariates which will yield the same predictions as M. The fact that $h_M^\varepsilon(\beta_M^\star) = 0$ for all subsets M with $|M| > n$ is a desirable consequence of defining M_T as a *least complex* sparse subset of covariates. It greatly reduces the number of candidate subsets of covariates to consider which helps to make the variable selection problem feasible, and with a provable consistency result. Reducing the number of candidate subsets is the basic motivation for the nonlocal priors of Johnson and Rossell (2012). The conditions necessary for the main result of this paper are presented next.

3.1. *Discussion of the conditions.* The first two conditions, Condition 3.1 and Condition 3.2, are to ensure that the true model, M_T, is identifiable. Recall that $\varepsilon_{M_T}(n,p)$ defines whether M_T is *least complex*. However, if $\varepsilon_{M_T}(n,p)$ is too large, then $h_{M_T}^\varepsilon(\beta_{M_T}^\star)$ will incorrectly decide that M_T is not a *least complex* subset of covariates. Condition 3.1 specifies how large $\varepsilon_{M_T}(n,p)$ can be whilst the true model remains identifiable. This condition turns out to be critically important in actual data applications.

CONDITION 3.1. *For large n and p,*

$$\varepsilon_{M_T}(n,p) \leq \frac{1}{18}\|X^T(\mu_T - Xb_{min})\|_2^2$$

where b_{min} solves

$$\min_{b\in\mathbb{R}^p} \frac{1}{2}\|X^T(\mu_T - Xb)\|_2^2 \quad \textit{subject to} \quad \|b\|_0 \leq p_T - 1.$$

Condition 3.2 is born from Lemma B.1 which is an important necessary result for the main result of this paper, Theorem 3.9. It is a slightly weaker version of condition (11) in Lai, Hannig and Lee (2015). They relate Condition 3.2 to the sparse Riesz condition (Zhang and Huang, 2008) which requires that the eigenvalues of $X_M^T X_M/n$ are uniformly bounded away from 0 and ∞. The requirements of this condition are reasonable because Δ_M grows very fast for M such that $M_T \not\subseteq M$. In fact, if M is orthogonal to M_T, then $\Delta_M = \|\mu_T\|_2^2$ which will be much larger than the denominator.

CONDITION 3.2. *For any $\gamma \geq 0$ where $p_T \leq \log(n)^\gamma$, and any $\alpha \in (0,1)$ where n^α represents the largest subset size considered,*

$$\lim_{\substack{n\to\infty \\ p\to\infty}} \min\left\{\frac{\Delta_M}{p_T \log(p)} : M \neq M_T, |M| \leq p_T\right\} = \infty,$$

$$\log(p) < \frac{n - p_T - 1}{4\log(n)^\gamma}, \quad \textit{and} \quad \liminf_{\substack{n\to\infty \\ p\to\infty}} \frac{n^{1-\alpha}}{\log(p)} > 2,$$

where $\Delta_M := \|\mu_T - H_M \mu_T\|_2^2$ as in Lai, Hannig and Lee (2015).

The term $\log(n)^\gamma$ in Condition 3.2 explicitly represents the sparsity assumption on the true model. γ is intended to serve as a tuning parameter for application on real data. In practice, $\gamma > 0$ should be chosen so that $\log(n)^\gamma$ is smaller than n, and can be used to reflect the practitioner's belief about the number of covariates in the true model. It will be shown empirically and on a real data set in Sections 4 and 5 that the variable selection procedure proposed in this paper is reasonably robust to the choice of γ. A penalized procedure such as LASSO or the Dantzig Selector also necessitates the choice of a tuning parameter, but in those cases the tuning parameter implies an assumption about the size of the nonzero coefficient estimates. Since γ explicitly reflects a belief about the number of covariates in the true model it is arguably more interpretable as a tuning parameter for a variable selection problem.

Conditions 3.3 and 3.4 demonstrate how large $\varepsilon_M(n,p)$ needs to be to achieve the consistency result of the main theorem. Condition 3.3 states that for large subsets of covariates $\varepsilon_M(n,p)$ needs to be larger than the *difference* in projections of the true model prediction, μ_T, onto M and onto a strict subset of M. This condition facilitates the intuition that the variable selection procedure will not concentrate on subsets M with redundant covariates. If there is a redundancy, then the difference in projections will be small so that the condition is easily achieved.

CONDITION 3.3. *For any M with $|M| > p_T$, for large n or p,*

$$\frac{9}{2}\|X^T(H_M - H_{M(-1)})\mu_T\|_2^2 < \varepsilon_M(n,p),$$

where $H_M := X_M(X_M^T X_M)^{-1} X_M^T$ is the projection matrix for subset M, and $H_{M(-1)}$ is the projection matrix for M after omitting the covariate which minimizes $\|X^T(H_M - H_{M(-1)})\mu_T\|_2^2$.

In fact, if $M_T \subset M$ with $p_T < |M|$, then the projection $H_M \mu_T = H_{M(-1)}\mu_T$ in which case Condition 3.3 holds trivially.

Lastly, Condition 3.4 describes the rate at which $\varepsilon_M(n,p)$ needs to grow to achieve the consistency of the main result. The distinction between Condition 3.3 and Condition 3.4 is that Condition 3.3 is a necessary lower bound for arguing that $\mathrm{E}[h^{\varepsilon}_M(\beta^{\star}_M)]$ vanishes for M such that $|M| > p_T$, while Condition 3.4 is the rate at which $\varepsilon_M(n,p)$ must grow to achieve the consistency result of Theorem 3.9. The terms which compete with $\varepsilon_M(n,p)$ arise in the proofs of Lemma B.1 and Theorem 3.9. Note that $\widehat{\sigma}^2_M := RSS_M/(n-p_M)$, where RSS_M is the classical residual sum of squares for model M.

CONDITION 3.4. *Defining* $\Lambda_M := tr\big((H_M X)^T H_M X\big)$,

$$\lim_{\substack{n\to\infty \\ p\to\infty}} \min_{|M|\le n^{\alpha}} \frac{\frac{\varepsilon_M(n,p)}{18\Lambda_M \widehat{\sigma}^2_M} + D_1 p_T - \varphi(M,n,p)}{\log(n)} = \infty$$

for any fixed constant $\alpha < 1$, *where*

$$\varphi(M,n,p) := 4e^2 n^{\alpha} + (D_1 + (1+4e^2)\log(p))|M| + \log\Big(\frac{|M|}{1-\frac{1}{n-|M|}}\Big),$$

and

$$D_1 = \frac{1}{2}\log\Big(\frac{6\pi}{1-\frac{n^{\alpha}+2}{n}}\Big).$$

Additionally, $\frac{\varepsilon_M(n,p)}{9\Lambda_M\widehat{\sigma}^2_M} < \frac{n-p_M}{2}$ *for all* M *with* $|M| \le n^{\alpha}$.

The term $\Lambda_M := \mathrm{tr}\big((H_M X)^T H_M X\big)$ arises from Lemmas 3.5 and 3.6. It is intimately related to the presence of collinearity amongst the covariates, and Condition 3.4 implies that $\varepsilon_M(n,p)$ must account for collinearity by controlling for Λ_M. Observe that if X is orthogonal, then $\Lambda_M = |M|$.

3.2. *Main result.* The first two results are lemmas which are needed in the proofs of Theorems 3.7 and 3.8. Lemma 3.5 illustrates the rate at which $\beta^{\star}_M$ concentrates around it's mean, $\widehat{\beta}_M$, the least squares estimator, and Lemma 3.6 illustrates the rate at which $\widehat{\beta}_M$ concentrates around it's mean, $\mathrm{E}(\widehat{\beta}_M)$. Theorem 3.7 uses these two lemmas to bound the rate at which $\beta^{\star}_M$ concentrates around $\mathrm{E}(\widehat{\beta}_M)$ for subsets M with $|M| > p_T$. This yields an upper bound on $\mathrm{E}[h^{\varepsilon}_M(\beta^{\star}_M)]$ with a probabilistic guarantee, and implies that $\mathrm{E}[h^{\varepsilon}_M(\beta^{\star}_M)]$ vanishes for large n and p, for large subsets M with redundancies. In what follows, recall that $\widehat{\sigma}^2_M := RSS_M/(n-p_M)$, where RSS_M is the

residual sum of squares for model M. The proofs are relegated to Appendix B.

LEMMA 3.5. *For any fixed positive constant $c_1 < 1$ assume $p_M \leq c_1 n$, and choose n and p such that $\frac{\varepsilon_M(n,p)}{9\Lambda_M \widehat{\sigma}_M^2} < \frac{n-p_M}{2}$. If*

$$\beta_M^\star \sim t_{n-p_M}\Big(\widehat{\beta}_M, \frac{RSS_M}{n-p_M}(X_M^T X_M)^{-1}\Big),$$

where $\widehat{\beta}_M = (X_M^T X_M)^{-1} X_M^T y$, then

$$P\Big(\frac{1}{2}\|X^T X_M(\beta_M^\star - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon_M(n,p)}{9}\Big) \leq \frac{2^{\frac{3}{2}} 3 p_M \widehat{\sigma}_M \sqrt{\Lambda_M} e^{-\frac{\varepsilon_M(n,p)}{18\Lambda_M \widehat{\sigma}_M^2}}}{\sqrt{\pi \varepsilon_M(n,p)}(1 - \frac{1}{n-p_M})}.$$

In the next lemma, P_{y} is used to denote the probability measure associated with the sampling distribution of the data y.

LEMMA 3.6. *Assume $p_M < n$, and $y \sim N(X_{M_T}\beta_{M_T}^\star, \sigma_M^{\star 2} I_n)$. Then the classical least squares estimator $\widehat{\beta}_M \sim N(E(\widehat{\beta}_M), \sigma_M^{\star 2}(X_M^T X_M)^{-1})$, and*

$$P_y\Big(\frac{1}{2}\|X^T X_M\big(\widehat{\beta}_M - E(\widehat{\beta}_M)\big)\|_2^2 \geq \frac{\varepsilon_M(n,p)}{9}\Big) \leq \frac{3 p_M \sigma_M^\star \sqrt{\Lambda_M}}{\sqrt{\pi \varepsilon_M(n,p)}} e^{-\frac{\varepsilon_M(n,p)}{9\sigma_M^{\star 2}\Lambda_M}}.$$

We continue by presenting a non-asymptotic concentration result for models that are larger than the true model.

THEOREM 3.7. *For any fixed positive constant $c_1 < 1$ suppose $p_T < p_M \leq c_1 n$, choose n and p such that $\frac{\varepsilon_M(n,p)}{9\Lambda_M \widehat{\sigma}_M^2} < \frac{n-p_M}{2}$, and assume that $\varepsilon_M(n,p)$ satisfies Condition 3.3. Then*

$$P_y\left(E[h_M^\varepsilon(\beta_M^\star)] \leq \frac{2^{\frac{3}{2}} 3 p_M \widehat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi \varepsilon_M(n,p)}(1 - \frac{1}{n-p_M})} e^{-\frac{\varepsilon_M(n,p)}{18\Lambda_M \widehat{\sigma}_M^2}}\right) \geq 1 - \frac{3 p_M \sigma_M^\star \sqrt{\Lambda_M}}{\sqrt{\varepsilon_M(n,p)\pi}} e^{-\frac{\varepsilon_M(n,p)}{9\sigma_M^{\star 2}\Lambda_M}}.$$

The next result is a probabilistic guarantee that if $\varepsilon_{M_T}(n,p)$ satisfies Condition 3.1, then the true model, M_T, will be a *least complex* subset of covariates. This result is a statement that M_T is identifiable.

THEOREM 3.8. *For any fixed positive constant $c_1 < 1$ suppose $p_T \leq c_1 n$, choose n and p such that $\frac{\varepsilon_{M_T}(n,p)}{9\Lambda_{M_T}\widehat{\sigma}^2_{M_T}} < \frac{n-p_T}{2}$, and assume that $\varepsilon_{M_T}(n,p)$ satisfies Condition 3.1. Then*

$$P_y\left(E[h^{\varepsilon}_{M_T}(\beta^{\star}_{M_T})] \geq 1 - \frac{2^{\frac{3}{2}}3p_{M_T}\widehat{\sigma}_{M_T}\sqrt{\Lambda_{M_T}}}{\sqrt{\pi\varepsilon_{M_T}(n,p)}(1-\frac{1}{n-p_{M_T}})}e^{-\frac{\varepsilon_{M_T}(n,p)}{18\Lambda_{M_T}\widehat{\sigma}^2_{M_T}}}\right) \geq 1 - \frac{3p_{M_T}\sigma^{\star}_M\sqrt{\Lambda_{M_T}}}{\sqrt{\varepsilon_{M_T}(n,p)\pi}}e^{-\frac{\varepsilon_{M_T}(n,p)}{9\sigma^{\star 2}_M\Lambda_{M_T}}}.$$

The following result is the main result of the paper. It shows that, for any fixed positive constant $\alpha < 1$, the ratio of the generalized fiducial probability of the true model to the sum over that of all other subsets of covariates M satisfying $|M| \leq n^{\alpha}$ will converge to 1 in probability for large n and p. Note that the restriction to subsets M with $|M| \leq n^{\alpha}$ is a stronger restriction than $|M| \leq c_1 n$ for any positive $c_1 < 1$, which was sufficient for Theorems 3.7 and 3.8. The reason being that the main result is much stronger than the results of these two theorems, and so it happens that only a smaller class of subsets can be considered. In fact, Theorems 3.7 and 3.8 are non-asymptotic results that hold for each model M separetaly, while Theorem 3.9 is an asymptotic result which applies uniformly over all M for a set of y vectors which converges to 1.

In a sparse setting it is assumed that $p_T << n$ so neglecting subsets with more than n^{α} covariates (for $\alpha < 1$) is reasonable. Recall that if $|M| > n$, then the observed covariates in M will necessarily be linearly dependent which implies that $h^{\varepsilon}_M(\beta^{\star}_M) = 0$, by definition.

THEOREM 3.9. *Suppose the true model is given by M_T, $\gamma > 0$ is some tuning parameter such that $p_T \leq \log(n)^{\gamma}$ for large n and p, and assume that Conditions 3.1-3.4 are satisfied. Then for any fixed positive constant $\alpha < 1$,*

$$r(M_T|y) = \frac{r(M_T|y)}{\sum_{j=1}^{n^{\alpha}}\sum_{M:|M|=j} r(M|y)} \xrightarrow{P} 1$$

as $n \to \infty$ or $n, p \to \infty$.

Although this is an asymptotic result, many of the ingredients that are used in its proof are non-asymptotic concentration results which are valid if Conditions 3.1-3.4 are satisfied. Therefore, it can be expected that when these conditions are satisfied the generalized fiducial distribution will concentrate on the true model M_T. This expectation is indeed validated by the empirical performance of the procedure, which is demonstrated in the remaining sections of the paper.

4. Simulation results. This section serves to demonstrate various aspects of the empirical performance of Algorithm A.2 on synthetic data. The focal point of the simulation on synthetic data is in the setting proposed in Ročková and George (2016) with $n = 100$, $p = 1000$, and only eight explanatory variables having nonzero coefficients. Following their setup the nonzero coefficients are chosen to be $(-2.5, -2, -1.5, -1, 1, 1.5, 2, 2.5)/\sqrt{3}$, and the model variance is set at 1 which makes the variable selection task considerably difficult. Two types of design matrices, X, are considered; one is generated with rows from the $\text{MVN}(0_p, I_p)$ distribution, and the other is generated with rows from the $\text{MVN}(0_p, \Sigma)$ distribution where the covariance matrix Σ has 1 in the diagonal components and 0.6 in the off-diagonal components. The latter design serves to impose a highly collinear structure among the explanatory variables.

Recall that a *least complex* subset of covariates will not be highly collinear. Thus, the highly collinear simulation setup of Ročková and George (2016) serves to evaluate the robustness of Algorithm A.2 to settings when the true model exhibits a high level of collinearity. This is a particularly difficult setting because a theoretical true subset of covariates is not well defined when it's covariates exhibit a high level of collinearity, and/or when it is highly correlated with other subsets of covariates.

Before discussing the results of our simulation study it is worth pointing out that Ročková and George (2016) only generate one data set (in the correlated design setting), and devote their empirical analysis to analyzing the solution paths of coefficient estimates for a sequence of penalty parameters. As a result, the empirical results of Ročková and George (2016) are not readily comparable to the results that follow, except to note that their variable selection procedure does not appear to be able to identify more than 6 of the true covariates. It seems that even for appropriately chosen penalty parameters, they are not able to retain the 2 covariates with the smallest signals. Our variable selection procedure has similar difficulties with finding these covariates, selecting them just shy of 20 percent of the time. However, below we use our theoretical results to determine how the simulation setup

must be modified in order to recover all 8 true covariates with our procedure.

The simulation here consists of generating 1,000 data sets (y and X pairs) in the independent design setup and 1,000 data sets in the correlated design setup. Algorithm A.2 is run for 100,000 steps, and the first 90,000 steps are discarded in analyzing the results.

Label descriptions for Tables 2-6

Avg Size	The number of covariates in each model visited for a given run of the MCMC algorithm, weighted by the proportion of times visited, and then averaged over all 1000 runs of the algorithm.
Avg PrRatio	For each of 1000 runs of the MCMC algorithm, with a test set of 5 data points, first compute the MSPE using the posterior mean vector and divide it by the MSPE of the true model using the true coefficient vector. Lastly, average this ratio over all 1000 runs of the MCMC algorithm. The posterior mean vector is computed as the weighted average of the posterior mean vectors from all subsets of covariates visited for a given run of the MCMC algorithm.
Avg FNR	The number of covariates incorrectly *excluded* from the most frequently occurring model in each of 1000 runs of the MCMC algorithm, then averaged over all 1000 values.
Avg FPR	The number of covariates incorrectly *included* in the most frequently occurring model in each of 1000 runs of the MCMC algorithm, then averaged over all 1000 values.
Avg CMR	The proportion of 1000 runs of the MCMC algorithm in which the most frequently occurring model was the correct model.
Avg COMR	The proportion of 1000 runs of the MCMC algorithm in which the most frequently occurring model contained all of the true covariates, and possibly others.
Incls Prop	For each of 1000 runs of the MCMC algorithm, the proportion of subsets visited which included a specified covariate, then averaged over all 1000 proportions.
Leng CI	For each of 1000 runs of the MCMC algorithm, first compute the weighted average empirical posterior CDF of a specified coefficient over all subsets in which the corresponding covariate was included. Then compute the length of the symmetric 95 percent empirical credible set. Lastly, average over all credible set length values.
Coverage	Compute the proportion of times that a specified true coefficient value lies within the credible set computed for 'Leng CI'.

TABLE 1

The choice of $\varepsilon_M(n,p)$ as suggested by Condition 3.4 is chosen to be

$$\varepsilon_M(n,p) = \Lambda_M \widehat{\sigma}_M^2 \Big(\frac{n^{0.51}}{9} + |M| \frac{\log(p\pi)^{1.1}}{9} - \log(n)^{\gamma} \Big)_+ \tag{4}$$

with $\gamma = 1.45$, and $N = 5$ sampled multivariate t vectors are taken at each proposed MCMC step (as is described in (1) of Algorithm A.2). The coefficients on the n, p, and $|M|$ terms were chosen ad hoc to make $\varepsilon_M(n,p)$ reasonably scaled for real data applications.

To initiate each simulated run of Algorithm A.2, a separate simple linear regression of y on each of the p covariates was performed, and the algorithm was started at the model consisting of the three covariates with the largest squared coefficient estimates. Then Algorithm A.1 was implemented iteratively on this initial subset of three covariates to move in the direction of *less complex* subsets until a *least complex* subset is obtained. The resulting subset is taken as the warm start to the algorithm. The reason for beginning with a very small subset of covariates is so that the algorithm will not get stuck at too large of a subset if γ is incorrectly assumed to be too big. Figures 1 and 2 and Tables 2 and 3 display the results of the two simulation setups. A detailed description of the labels for Tables 2-6 is given in Table 1.

MVN$(0_p, I_p)$ **design matrix rows with** $n = 100$, $p = 1000$, **and** $\gamma = 1.45$

Average	Size	std err	PrRatio	std err	FNR	FPR	CMR	COMR
	8.3756	0.4715	1.3608	1.0810	0.1890	0.5680	0.4660	0.8290
	β_1^0	β_2^0	β_3^0	β_4^0	β_5^0	β_6^0	β_7^0	β_8^0
True Coef	-1.443	-1.154	-0.866	-0.577	0.577	0.866	1.154	1.443
Incls Prop	1.0000	0.9999	0.9989	0.9145	0.8909	0.9964	0.9999	0.9990
Leng CI	0.6716	0.5401	0.4520	0.4143	0.4143	0.4556	0.5551	0.6832
Coverage	0.9740	0.9630	0.9349	0.9463	0.9504	0.9489	0.9530	0.9750

TABLE 2

Simulation consists of 1000 runs of Algorithm A.2. 100,000 MCMC steps, burn-in at 90,000 steps. The ratio for Condition 3.1 is $0.2216 < 0.5$.

MVN$(0_p, \Sigma)$ **design matrix rows with** $n = 100$, $p = 1000$, **and** $\gamma = 1.45$

Average	Size	std err	PrRatio	std err	FNR	FPR	CMR	COMR
	6.9995	0.0003	1.9187	1.7013	2.1090	1.1090	0.0000	0.0000
	β_1^0	β_2^0	β_3^0	β_4^0	β_5^0	β_6^0	β_7^0	β_8^0
True Coef	-1.443	-1.154	-0.866	-0.577	0.577	0.866	1.154	1.443
Incls Prop	0.9914	0.9505	0.7002	0.1853	0.1792	0.7028	0.9527	0.9931
Leng CI	0.7921	0.7627	0.7224	0.6980	0.7066	0.7204	0.7593	0.7890
Coverage	0.9359	0.9304	0.9279	0.9231	0.9161	0.9224	0.9470	0.9419

TABLE 3

Simulation consists of 1000 runs of Algorithm A.2. 100,000 MCMC steps, burn-in at 90,000 steps. Σ *has the value 1 in the diagonal components and 0.6 in the off-diagonal components. The ratio for Condition 3.1 is* $6.3841 > 0.5$.

MVN$(0_p, \Sigma)$ **design matrix rows with** $n = 100$**,** $p = 500$**, and** $\gamma = 1.45$

Average	Size	std err	PrRatio	std err	FNR	FPR	CMR	COMR
	7.9991	0.0005	1.6606	1.3417	1.4400	1.4400	0.1970	0.1970
	β_1^0	β_2^0	β_3^0	β_4^0	β_5^0	β_6^0	β_7^0	β_8^0
True Coef	-1.443	-1.154	-0.866	-0.577	0.577	0.866	1.154	1.443
Incls Prop	0.9979	0.9775	0.8269	0.3364	0.3467	0.8201	0.9848	0.9998
Leng CI	0.7484	0.7350	0.7033	0.6778	0.6788	0.7093	0.7398	0.7560
Coverage	0.9429	0.9396	0.9303	0.9356	0.9296	0.9328	0.9438	0.9440

TABLE 4

Simulation consists of 1000 runs of Algorithm A.2. 100,000 MCMC steps, burn-in at 90,000 steps. Σ *has the value 1 in the diagonal components and 0.6 in the off-diagonal components. The ratio for Condition 3.1 is* $0.0 < 0.5$.

MVN$(0_p, I_p)$ **design matrix rows with** $n = 500$**,** $p = 1000$**, and** $\gamma = 0.5$

Average	Size	std err	PrRatio	std err	FNR	FPR	CMR	COMR
	7.9915	0.0891	1.0328	0.1841	0.0080	0.0000	0.9920	0.9920
	β_1^0	β_2^0	β_3^0	β_4^0	β_5^0	β_6^0	β_7^0	β_8^0
True Coef	-1.443	-1.154	-0.866	-0.577	0.577	0.866	1.154	1.443
Incls Prop	1.0000	1.0000	0.9999	0.9939	0.9979	0.9999	0.9999	1.0000
Leng CI	0.3309	0.2546	0.1808	0.1770	0.1771	0.1809	0.2513	0.3296
Coverage	0.9760	0.9600	0.9460	0.9447	0.9649	0.9690	0.9720	0.9870

TABLE 5

Simulation consists of 1000 runs of Algorithm A.2. 30,000 MCMC steps, burn-in at 20,000 steps. The ratio for Condition 3.1 is $0.4426 < 0.5$.

MVN$(0_p, \Sigma)$ **design matrix rows with** $n = 20,000$**,** $p = 30$**, and** $\gamma = 0.5$

Average	Size	std err	PrRatio	std err	FNR	FPR	CMR	COMR
	7.9909	0.0631	1.0019	0.0265	0.0020	0.0000	0.9990	0.9990
	β_1^0	β_2^0	β_3^0	β_4^0	β_5^0	β_6^0	β_7^0	β_8^0
True Coef	-1.443	-1.154	-0.866	-0.577	0.577	0.866	1.154	1.443
Incls Prop	1.0000	0.9992	0.9986	0.9976	0.9976	0.9986	0.9993	1.0000
Leng CI	0.3573	0.1918	0.0421	0.0420	0.0421	0.0421	0.2081	0.3575
Coverage	0.9580	0.9600	0.9570	0.9560	0.9499	0.9460	0.9540	0.9550

TABLE 6

Simulation consists of 1000 runs of Algorithm A.2. 1,000 MCMC steps, burn-in at 300 steps. Σ *has the value 1 in the diagonal components and 0.6 in the off-diagonal components. The ratio for Condition 3.1 is* $0.3134 < 0.5$.

In the independent design setup, Table 2 shows that the generalized fiducial distribution concentrates, on average, on the true number covariates. The average MSPE is only about 36 percent larger than that of the true model, and is likely because the two covariates corresponding to the smallest (in magnitude) coefficients were never included in about 10 percent of the simulated runs. The true subset of covariates is the most frequently occurring subset in about 47 percent of the runs, and a subset containing the truth

is the most frequently occurring subset in about 83 percent of the runs. The results are not as good in the correlated design setup. In fact, it is observed from Table 3 that the MCMC algorithm never selects the true model nor a model containing the truth, but rather it concentrates on subsets contained within the true model (as can be seen from the Avg FNR and Avg FPR). However, the lack in performance seen in Table 2 and particularly in Table 3 do not illustrate deficiencies in the variable selection procedure, but rather they illustrate the importance of Conditions 3.1-3.4.

Particularly important are Conditions 3.1 and 3.2 which are violated in these simulation setups. Computing the ratio of terms in the statement of Condition 3.1, $\varepsilon_{M_T}(n,p)/\|X^T(\mu_T - Xb_{min})\|_2^2$, and averaging over all 1000 data sets yields a value of approximately 6.38 for the correlated design setting, and a value of about 0.22 for the independent design setting. In practical applications, Condition 3.1 was found to be violated when the ratio exceeded 1/2, as opposed to the 1/18. The factor of 1/2 is reasonable as it is found in the definition of h_M^{ε} in (1). This condition is clearly violated in the collinear setting, and explains why the average model size is less than the true number of covariates for the collinear setup but not for the independent setup.

To more fully appreciate the importance of Condition 3.1, the setup in the collinear setting is modified so that this condition is satisfied with the same assumption about the true model. Accordingly, γ and n are held fixed at 1.45 and 100, respectively since $p_T \leq \log(n)^{\gamma}$. It is determined that Condition 3.1 is satisfied in the collinear setting when p is reduced to 500 covariates. The results of this revised setup shown Table 4 are markedly improved. However, the two covariates with the smallest magnitude coefficients are only included in about 90 percent of the visited subsets in the independent setting, and in about 35 percent in the collinear setting. This is due to the violation of Condition 3.2.

Specifically, Condition 3.2 is violated because Δ_M is not large enough for subsets M with $|M| \leq p_T$, but this is hard to explicitly verify. The meaning of this is that the signal-to-noise ratio is too small for the covariates with coefficient magnitudes of 0.577, a value which is far from the tails of the standard normal distribution. The second stipulation in Condition 3.2 is also violated, since

$$\frac{4\log(100)^{1.45}\log(500)}{100-8-1} = 2.50 > 1.$$

The violation of this stipulation implies that the residual sum of squares cannot be controlled for subsets M with $|M| \leq p_T$ (see the proof of Lemma

B.1 for the details). Note that the coefficient of 4 can be replaced by any number greater than 3. In essence, the violation of Condition 3.2, as a whole, means that p is too large relative to the size of n.

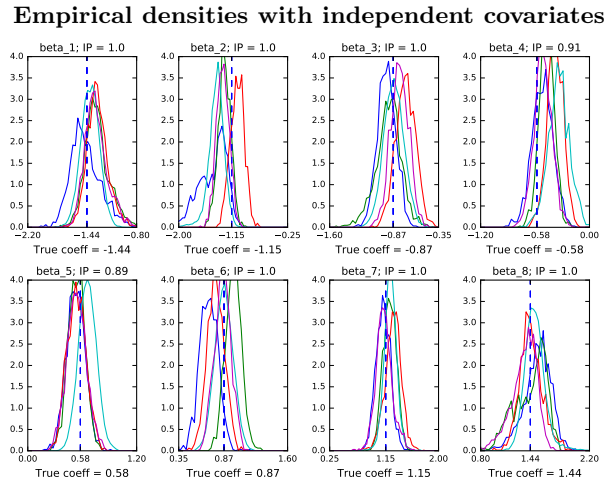
Empirical densities with independent covariates
beta_1; IP = 1.0
beta_2; IP = 1.0
beta_3; IP = 1.0
beta_4; IP = 0.91
True coeff = -1.44
True coeff = -1.15
True coeff = -0.87
True coeff = -0.58
beta_5; IP = 0.89
beta_6; IP = 1.0
beta_7; IP = 1.0
beta_8; IP = 1.0
True coeff = 0.58
True coeff = 0.87
True coeff = 1.15
True coeff = 1.44

FIG 1. *Empirical density curves for each of the true coefficients, for 5 runs of the algorithm in the $MVN(0_p, I_p)$ design setup.*

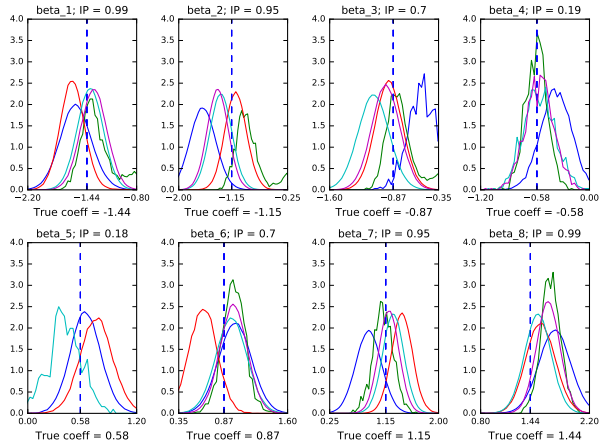
Empirical densities with correlated covariates
beta_1; IP = 0.99
beta_2; IP = 0.95
beta_3; IP = 0.7
beta_4; IP = 0.19
True coeff = -1.44
True coeff = -1.15
True coeff = -0.87
True coeff = -0.58
beta_5; IP = 0.18
beta_6; IP = 0.7
beta_7; IP = 0.95
beta_8; IP = 0.99
True coeff = 0.58
True coeff = 0.87
True coeff = 1.15
True coeff = 1.44

FIG 2. *Empirical density curves for each of the true coefficients, for 5 runs of the algorithm in the $MVN(0_p, \Sigma)$ design setup. Σ has the value 1 in the diagonal components and 0.6 in the off-diagonal components.*

However, while Condition 3.2 is necessary asymptotically, it is less in-

formative in applications than Condition 3.1 because the latter is directly related to the size of $\varepsilon_M(n,p)$, which is actually computed in the computer program to run the algorithm.

To test the sensitivity of the variable selection procedure to the choice of γ choose $\gamma = 0.5$, a value which reflects a much smaller true model (since $\log(100)^{0.5} \approx 2$), and find n and p which ensure that all of Conditions 3.1-3.4 are satisfied. Accordingly, Table 5 shows the boundary of how large n must be for $p = 1000$ in the independent setting, and Table 6 shows the boundary sizes of n and p in the collinear setting. The results for the collinear setting require $n = 20,000$ and $p = 30$ which are extreme, but only necessary because all covariates have the very high correlation of 0.6, including those in the true model. It is observed in these results that when Conditions 3.1-3.4 hold, the choice of γ is of no consequence.

5. Real data. The purpose of this section is to evaluate the performance of Algorithm A.2 on a real data set. The data set chosen for analysis was also used in Bondell and Reich (2012) and Narisetty and He (2014). Narisetty and He (2014) provide the following description of the data set. It consists of 22,575 gene expressions from 31 female and 29 male inbred mice, and also includes the numbers of phenotypes for phosphoenopyruvate carboxykinase (PEPCK) and glycerol-3-phosphate acyltransferase (GPAT). More information as well as the data can be found at Gene Expression Omnibus (http://www.ncbi.nlm.nih.gov/geo; accession number GSE3330).

Following the analysis of Narisetty and He (2014) the 22,575 covariates were pre-screened based on the magnitudes of their marginal correlation with the response variables. Separate pre-screenings were carried out for both of the responses, PEPCK and GPAT. One setup consisted of $p = 200$ covariates, and the other consisted of $p = 400$ covariates (both cases included an intercept and a variable for gender). As done in Bondell and Reich (2012) and Narisetty and He (2014), the data was randomly split into a training set of 55 observations and a test set of 5 observations. Then Algorithm A.2, using $\varepsilon_M(n,p)$ from (4), was run on 100 such random permutations of the training and test sets to estimate the out-of-sample prediction error. The intercept and gender variable were supplied in each visited subset when computing the MSPE.

Figure 3 shows the results of the analysis in two plots which are comparable to those of Figure 1 in Narisetty and He (2014). The average MSPE is plotted against various choices of γ, and the average model sizes are in parentheses above each point. Average model sizes are computed by taking the weighted average of the sizes of the subsets visisted in the MCMC chain

for each of the 100 runs of the algorithm, and then averaging over the 100 runs of the algorithm. The reader will observe that the number of included covariates increases with γ. This is a consequence of the residual sum of squares term in the mass function (2) because the residual sum of squares always puts upward pressure on model complexity.

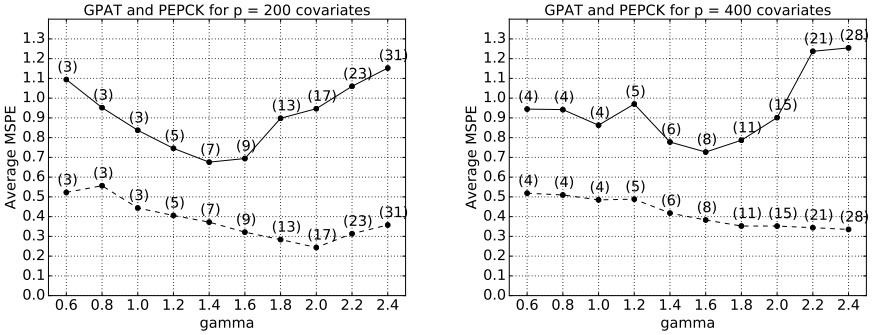

FIG 3. *Model sizes averaged over 100 permutations of the training and test set are in parentheses, rounded to the nearest integer. 55 data points are in the training set and 5 are in the test set. 30,000 MCMC steps, burn-in at 20,000 steps. The dashed lines correspond to PEPCK and the solid lines correspond to GPAT.*

In comparison to Narisetty and He (2014), Algorithm A.2 finds models with similar, but slightly larger MSPE. Note however that for this data set, for the majority of the γ values the second stipulation in Condition 3.2 does not hold, which hinders the performance of the algorithm (as discussed in Section 4). It is also worth noting that the results are not directly comparable to those of Narisetty and He (2014) because the variable selection method proposed here is foremost concerned with managing model complexity over marginal improvements in prediction error.

References.

ANDRIEU, C. and ROBERTS, G. (2009). The psuedo-marginal approach for efficient monte carlo computations. *The Annals of Statistics* **37** 697-725.

BEAUMONT, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics* **164** 1139-1160.

BERK, R. A. (2008). *Statistical Learning from a Regression Perspective*. Springer, New York, NY.

BERTSIMAS, D., KING, A. and MAZUMDER, R. (2016). Best Subset Selection via a Modern Optimization Lens. *The Annals of Statistics* **44** 813-852.

BONDELL, H. D. and REICH, B. J. (2012). Consistent high-dimensional Bayesian variable selection via penalized credible regions. *Journal of the American Statistical Association* **107** 1610-1624.

CANDES, E. and TAO, T. (2007). The Dantzig Selector: Statistical estimation when p is much greater than n. *The Annals of Statistics* **35** 2313-2351.

Fan, J. and Li, R. (2001). Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properities. *Journal of the American Statistical Association* **96** 1348-1360.
Ghosh, J. and Ghattas, A. E. (2015). Bayesian Variable Selection Under Collinearity. *The American Statistician* **69** 649-660.
Hannig, J., Iyer, H., Lai, R. C. S. and Lee, T. C. M. (2016). Generalized Fiducial Inference: A Reviev and New Results. *Journal of American Statistical Association* **111** 1346-1361.
Jameson, G. J. O. (2013). Inequalities for gamma function ratios. *American Math. Monthly* **120** 936-940.
Johnson, V. E. and Rossell, D. (2012). Bayesian Model Selection in High-Dimensional Settings. *Journal of the American Statistical Association* **107**.
Lai, R. C. S., Hannig, J. and Lee, T. C. M. (2015). Generalized Fiducial Inference for Ultrahigh-Dimensional Regression. *Journal of the American Statistical Association* **110** 760-772.
Lan, H., Chen, M., Flowers, J. B., Yandell, B. S., Stapleton, D. S., Mata, C. M., Mui, E. T., Flowers, M. T., Schueler, K. L., Manly, K. F., Williams, R. W., Kendziorski, K. and Attie, A. D. (2006). Combined expresssion trait correlations and expression quantitative trait locus mapping. *PLoS Genetics* **6e**.
Luo, S. and Chen, Z. (2013). Extended BIC for Linear Regression Models With Diverging Number of Relevant Features and High or Ultra-High Feature Spaces. *Journal of Statistical Planning and Inference* **143** 494-504.
Narisetty, N. N. and He, X. (2014). Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics* **24** 789-817.
Ročková, V. and George, E. I. (2016). The Spike-and-Slab LASSO. *Journal of the American Statistical Association.*
Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society* **58** 267-288.
Yuan, M. and Lin, Y. (2005). Efficient Empirical Bayes Variable Selection and Estimation in Linear Models. *Journal of the American Statistical Association* **100** 1215-1225.
Zhang, C. and Huang, J. (2008). The Sparsity and Bias of the Lasso Selection in High-Dimensional Linear Regression. *The Annals of Statistics* **36** 1567-1594.

APPENDIX A: TECHNICAL DETAILS FOR ALGORITHM COMPUTATIONS

A.1. Evaluating the model complexity decision function. The purpose of this section is to provide the technical details for evaluating h_M^ε as defined in (1). Algorithm A.1 which is adapted from Bertsimas, King and Mazumder (2016) is implemented for this purpose, and the details are as follows.

Evaluating $h_M^\varepsilon(\beta_M^\star)$ amounts to solving

$$\min_{b\in\mathbb{R}^p} g(b) \quad \text{subject to} \quad \|b\|_0 \le |M| - 1,$$

where

$$g(b) := \frac{1}{2}\|X^T(X_M\beta_M^\star - Xb)\|_2^2.$$

As discussed in Bertsimas, King and Mazumder (2016), this L_0 minimization problem can be solved with Algorithm A.1 since $g(b) \geq 0$ is convex and has Lipschitz continuous gradient:

$$\nabla g(b) = X^T X X^T (Xb - X_M \beta_M^\star) \quad \text{and}$$
$$\|\nabla g(b) - \nabla g(\widetilde{b})\|_2 \leq \lambda_{max}((X^T X)^2)\|b - \widetilde{b}\|_2,$$

where $\lambda_{max}((X^T X)^2)$ is the maximum of the eigenvalues of $(X^T X)^2$.

The basic intuition is to update the solution vector iteratively in a gradient decent fashion. The cardinality constraint is imposed by only retaining the $|M| - 1$ largest in magnitude vector components in the gradient direction, at every iteration.

ALGORITHM A.1. **(1)** *Initialize with some $b^{(0)} \in \mathbb{R}^p$ with $\|b^{(0)}\|_0 \leq |M|$, and set $b^{(1)} = b^{(0)}_{-1}$ where $b^{(0)}_{-1}$ is the vector $b^{(0)}$ with it's smallest component (in absolute value) removed.*

(2) *For $m \geq 1$, set*

$$b_i^{(m+1)} = \begin{cases} c_i & \text{if } \ i \in \{(1), \ldots, (|M| - 1)\} \\ 0 & \text{else} \end{cases}, \quad \text{for } \ i \in \{1, \ldots, p\},$$

where

$$c = b^{(m)} - \frac{1}{l}\nabla g(b^{(m)}) = b^{(m)} - \frac{X^T X X^T (Xb^{(m)} - X_M \beta_M^\star)}{\lambda_{max}((X^T X)^2)},$$

and $|c_{(1)}| \geq |c_{(2)}| \geq \cdots \geq |c_{(p)}|$.

(3) *Repeat until one of the following conditions are satisfied.*

(i) *$g(b^{(m+1)}) = \frac{1}{2}\|X^T(X_M \beta_M^\star - Xb^{(m+1)})\|_2^2 < \varepsilon_M(n, p)$, or*

(ii) *$g(b^{(m)}) - g(b^{(m+1)})$ is arbitrarily small (not in absolute value), or*

(iii) *Some maximum number of iterations has been exceeded.*

Bertsimas, King and Mazumder (2016) provide provable guarantees that the algorithm will converge to a 'first-order stationary point'. A first-order stationary point is a vector $\tilde{b} \in \mathbb{R}^p$ with $\|\tilde{b}\|_0 \leq |M| - 1$ which satisfies

$$\tilde{b} = \tilde{b} - \frac{1}{l}\nabla g(\tilde{b}).$$

Observe that if the distribution of one of the components of $\beta_M^\star$ is concentrated at zero, then their results guarantee $h_M^\varepsilon(\beta_M^\star) = 0$ by way of Algorithm

A.1. In this case, the problem is made even simpler because $X_M\beta^\star_M$ is noiseless and can be exactly recovered. Paraphrasing from Bertsimas, King and Mazumder (2016), they explain that Algorithm A.1 detects the active set after a few iterations, and then takes additional time to estimate the coefficient values to a high accuracy level. This translates here to fast computation of h^ε_M predominately because highly accurate coefficient estimates are not always necessary due to condition (i) of step (3). If this condition is met, then the algorithm has determined that $h^\varepsilon_M(\beta^\star_M) = 0$, and thus the subset of covariates M is not *least complex*.

A.2. Setting up the MCMC algorithm. This section serves to provide the details of GIMH from Andrieu and Roberts (2009) used to compute the subset probabilities, $r(M|y)$ as in (2). Begin by defining

$$r(M, v|y) := C \cdot \pi^{\frac{p_M}{2}} \Gamma\Big(\frac{n - p_M}{2}\Big) RSS_M^{-(\frac{n-p_M-1}{2})} h^\varepsilon_M(v)$$

which is not a probability density function. Note that C is the normalizing constant from (2). Further, let

$$r_M(v|y) := \frac{r(M, v|y)Q_M(v)}{\int r(M, v|y)Q_M(v)\ dv}$$

denote the conditional density of v given a subset of covariates M, where $Q_M(v)$ is the density function associated with the location scale T distribution (3). Then

$$\frac{r_M(v|y)}{Q_M(v)} \underbrace{\int r(M, v|y)Q_M(v)\ dv}_{= r(M|y)} = r(M, v|y). \tag{5}$$

Lastly, let the columns $B(i)$ of a new matrix B consist of a sample of size N from distribution (3), and denote the joint density function of the sample as $Q^N_M(B) := \prod_{i=1}^N Q_M(B(i))$, by independence. Then, in the convention of Andrieu and Roberts (2009), the GIMH algorithm has target distribution

$$\begin{aligned} r^N(M, B|y) &:= r(M|y) \cdot Q^N_M(B) \cdot \frac{1}{N} \sum_{i=1}^N \frac{r_M(B(i)|y)}{Q_M(B(i))} \\ &= Q^N_M(B) \cdot \frac{1}{N} \sum_{i=1}^N r(M, B(i)|y), \end{aligned}$$

where the second line is true by (5). Observe that $r^N(M, B|y)$ has the desired distribution, $r(M|y)$, as its marginal distribution. The results of Andrieu and

Roberts (2009) guarantee that MCMC with target distribution $r^N(M, B|y)$ will produce samples of M according to $r(M|y)$ asymptotically, as long as N is large enough.

Use $M_{(t)}$ and $B_{(t)}$ to denote the subset of covariates and sample of vectors, respectively, at step t of the GIMH algorithm. Then at step $t+1$ propose a new model, $\widetilde{M} \sim q(\cdot|M_{(t)})$, and a new sample of vectors, $\widetilde{B} \sim Q^N_{\widetilde{M}}(\cdot)$. This results in the following acceptance ratio

$$\begin{aligned}\rho(M_{(t)}, \widetilde{M}) &= \min\left\{\frac{r^N(\widetilde{M}, \widetilde{B}|y)q(M_{(t)}|\widetilde{M})Q^N_{M_{(t)}}(B_{(t)})}{r^N(M_{(t)}, B_{(t)}|y)q(\widetilde{M}|M_{(t)})Q^N_{\widetilde{M}}(\widetilde{B})}, 1\right\} \\ &= \min\left\{\frac{\left[\frac{1}{N}\sum_{i=1}^N r(\widetilde{M}, \widetilde{B}(i)|y)\right]q(M_{(t)}|\widetilde{M})}{\left[\frac{1}{N}\sum_{i=1}^N r(M_{(t)}, B_{(t)}(i)|y)\right]q(\widetilde{M}|M_{(t)})}, 1\right\}.\end{aligned} \tag{6}$$

The pseudo-code for the constructed MCMC algorithm is presented next.

ALGORITHM A.2. *Given some subset, $M_{(t)}$, of the p covariates at time t,*

(1) Sample.

$$\widetilde{M} = \begin{cases} M_{(t)} \cup \{a \text{ new covariate}\} & w.p. \ \frac{1}{3} \\ M_{(t)} \setminus \{an \text{ existing covariate}\} & w.p. \ \frac{1}{3} \\ \big(M_{(t)} \setminus \{an \text{ existing covariate}\}\big) \cup \{a \text{ new covariate}\} & w.p. \ \frac{1}{3} \end{cases}$$

where a covariate is added to the subset $M_{(t)}$ with probability $w_j^{(t)}$ for $j \in \{1, \ldots, p - |M_{(t)}|\}$, and is dropped from $M_{(t)}$ with probability $v_i^{(t)}$ for $i \in \{1, \ldots, |M_{(t)}|\}$. This yields the proposal probability function

$$q(\widetilde{M}|M_{(t)}) = \begin{cases} \frac{1}{3}w_j^{(t)} & if \ \ |\widetilde{M}| > |M_{(t)}| \\ \frac{1}{3}v_i^{(t)} & if \ \ |\widetilde{M}| < |M_{(t)}| \\ \frac{1}{3}w_j^{(t)}v_i^{(t)} & if \ \ |\widetilde{M}| = |M_{(t)}| \end{cases},$$

for $j \in \{1, \ldots, p - |M_{(t)}|\}$ and $i \in \{1, \ldots, |M_{(t)}|\}$. The vectors $\vec{w}^{(t)}$ and $\vec{v}^{(t)}$ are vectors of weights depending on $M_{(t)}$, which sum to 1.

Given the proposal $\widetilde{M}$, for $k \in \{1, \ldots, N\}$ generate

$$\widetilde{B}(k) \sim t_{n-|\widetilde{M}|}\Big((X^T_{\widetilde{M}}X_{\widetilde{M}})^{-1}X^T_{\widetilde{M}}y, \frac{RSS_{\widetilde{M}}}{n - |\widetilde{M}|}(X^T_{\widetilde{M}}X_{\widetilde{M}})^{-1}\Big).$$

(2) Update.

$$M_{(t+1)} = \begin{cases} \widetilde{M} & w.p. \;\; \rho(M_{(t)}, \widetilde{M}) \\ M_{(t)} & w.p. \;\; 1 - \rho(M_{(t)}, \widetilde{M}) \end{cases}$$

where the acceptance ratio is given by $\rho(M_{(t)}, \widetilde{M})$ *as in (6).*

A natural choice of weights are

$$w_j^{(t)} := \frac{\widehat{\beta}_j^2}{\sum_{k=1}^{p-|M_{(t)}|} \widehat{\beta}_k^2}, \quad \text{for} \;\; j \in \{1, \ldots, p - |M_{(t)}|\},$$

and

$$v_i^{(t)} = \frac{\widehat{\beta}_i^{-2}}{\sum_{k=1}^{|M_{(t)}|} \widehat{\beta}_k^{-2}}, \quad \text{for} \;\; i \in \{1, \ldots, |M_{(t)}|\},$$

where the coefficient estimates are the least squares estimates for the simple linear regression of each covariate on the response, y, separately. This allows for reasonable estimates of the weights, even in the face of a highly collinear system or when $p > n$. Empirical performance on synthetic data suggests that using these weights greatly improves the time it takes for the algorithm to propose the true covariates.

APPENDIX B: PROOFS

Proof of Lemma 3.5. From the distributional assumption on $\beta_M^\star$ it follows that

$$T := \sqrt{\frac{n - p_M}{RSS_M}} (X_M^T X_M)^{\frac{1}{2}} (\beta_M^\star - \widehat{\beta}_M) \sim \mathrm{t}_{n-p_M}(0, I_{p_M}).$$

Thus,

$$\|X^T X_M (\beta_M^\star - \widehat{\beta}_M)\|_2^2 = \|AT\|_2^2 \frac{RSS_M}{n - p_M} = \|UDW^T T\|_2^2 \frac{RSS_M}{n - p_M}$$

where $A = X^T X_M (X_M^T X_M)^{-\frac{1}{2}}$ is a $p \times p_M$ matrix which has the singular value decomposition $A = UDW^T$ for U and W each orthogonal matrices. By the definition of the multivariate t distribution,

$$T = \sqrt{\frac{n - p_M}{V}} Z \implies W^T T = \sqrt{\frac{n - p_M}{V}} W^T Z =: \sqrt{\frac{n - p_M}{V}} \widetilde{Z} =: \widetilde{T},$$

where $V \sim \chi^2_{n-p_M}$ and $Z, \widetilde{Z} \sim \mathrm{MVN}(0, I_{p_M})$. Then

$$\|AT\|_2^2 = \|UD\widetilde{T}\|_2^2 = \widetilde{T}^T D^T D \widetilde{T} = \sum_1^{p_M} \widetilde{T}_i^2 \lambda_i \leq \Lambda_M \max \widetilde{T}_j^2$$

where $\Lambda_M = \sum_1^{p_M} \lambda_i$, and λ_i is the ith eigenvalue of $A^T A$. Observe that Λ_M also has the more intuitive expression $\Lambda_M = \mathrm{tr}(A^T A) = \mathrm{tr}\big((H_M X)^T H_M X\big)$ which illustrates that it is intimately related to the presence of collinearity amongst the covariates.

For conciseness, denote $\varepsilon := \varepsilon_M(n,p)$, and recall that $\widehat{\sigma}_M^2 := RSS_M/(n - p_M)$. Then

$$\begin{aligned} P\Big(\frac{1}{2}\|X^T X_M(\beta_M^\star - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\Big) &= P\Big(\frac{1}{2}\|AT\|_2^2 \frac{RSS_M}{n-p_M} \geq \frac{\varepsilon}{9}\Big) \\ &\leq P\Big(\frac{1}{2}\Lambda_M \max \widetilde{T}_j^2 \frac{RSS_M}{n-p_M} \geq \frac{\varepsilon}{9}\Big) \\ &= P\Big(\frac{n-p_M}{V} \max Z_j^2 \widehat{\sigma}_M^2 \geq \frac{2\varepsilon}{9\Lambda_M}\Big). \end{aligned}$$

Since $V \sim \chi^2_{n-p_m}$,

$$\begin{aligned} &P\Big(\frac{1}{2}\|X^T X_M(\beta_M^\star - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\Big) \\ &\quad\leq \int_0^\infty P\Big(\max|Z_j| \geq \frac{\sqrt{2\varepsilon v/(n-p_M)}}{3\widehat{\sigma}_M\sqrt{\Lambda_M}}\Big) \frac{v^{\frac{n-p_M}{2}-1}e^{-\frac{v}{2}}}{2^{\frac{n-p_M}{2}}\Gamma(\frac{n-p_M}{2})}\,dv \\ &\quad\leq \int_0^\infty \sum_{j=1}^{p_M} P\Big(|Z_j| \geq \frac{\sqrt{2\varepsilon v/(n-p_M)}}{3\widehat{\sigma}_M\sqrt{\Lambda_M}}\Big) \frac{v^{\frac{n-p_M}{2}-1}e^{-\frac{v}{2}}}{2^{\frac{n-p_M}{2}}\Gamma(\frac{n-p_M}{2})}\,dv \\ &\quad\leq \frac{3p_M\widehat{\sigma}_M\sqrt{\Lambda_M(n-p_M)}}{\sqrt{2\pi\varepsilon}\left(1+\frac{\varepsilon}{9\Lambda_M\widehat{\sigma}_M^2(\frac{n-p_M}{2})}\right)^{\frac{n-p_M}{2}-\frac{3}{2}}} \cdot \frac{\Gamma(\frac{n-p_M-1}{2})}{\Gamma(\frac{n-p_M}{2})} \end{aligned} \tag{7}$$

where the last inequality follows because for the standard normal CDF, Φ for $x > 0$,

$$\Phi(-x) \leq \frac{\frac{1}{\sqrt{2\pi}}e^{\frac{-x^2}{2}}}{x}.$$

To simplify the bound, observe first that Jameson (2013) gives

$$\frac{\Gamma(\frac{n-p_M-1}{2})}{\Gamma(\frac{n-p_M}{2})} \leq \frac{\sqrt{2(n-p_M)}}{n-p_M-1}. \tag{8}$$

Second, observe that for $0 \leq x \leq n$,

$$\Big(1+\frac{x}{n}\Big)^{-n} \leq e^{-x+\frac{x^2}{2n}} = e^{-x(1-\frac{x}{2n})} \leq e^{-\frac{x}{2}}.$$

By assumption $p_M \leq c_1 n$ for any fixed positive constant $c_1 < 1$, and $\frac{\varepsilon_M(n,p)}{9\Lambda_M \widehat{\sigma}_M^2} < \frac{n-p_M}{2}$ which implies that

$$\left(1 + \frac{\varepsilon}{9\Lambda_M \widehat{\sigma}_M^2 (\frac{n-p_M}{2})}\right)^{-\frac{n-p_M}{2}+\frac{3}{2}} \leq 2^{\frac{3}{2}} e^{-\frac{\varepsilon}{18\Lambda_M \widehat{\sigma}_M^2}}. \tag{9}$$

Therefore, applying (8) and (9) to (7) gives

$$P\Big(\frac{1}{2}\|X^T X_M(\beta_M^\star - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\Big) \leq \frac{2^{\frac{3}{2}} 3 p_M \widehat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi\varepsilon}(1 - \frac{1}{n-p_M})} e^{-\frac{\varepsilon}{18\Lambda_M \widehat{\sigma}_M^2}}.$$

■

Proof of Lemma 3.6. Similar to the proof of Lemma 3.5. ■

Proof of Theorem 3.7. Recall that

$$h_M^\varepsilon(\beta_M^\star) = 1\Big\{\frac{1}{2}\|X^T(X_M\beta_M^\star - Xb_{min})\|_2^2 \geq \varepsilon_M(n,p)\Big\},$$

where b_{min} solves

$$\min_{b\in\mathbb{R}^p} \frac{1}{2}\|X^T(X_M\beta_M^\star - Xb)\|_2^2 \quad \text{subject to} \quad \|b\|_0 \leq |M| - 1.$$

For conciseness, denote $\varepsilon := \varepsilon_M(n,p)$. Observe that

$$\begin{aligned} \mathrm{E}[h_M^\varepsilon(\beta_M^\star)] &= P\Big(\frac{1}{2}\|X^T(X_M\beta_M^\star - Xb_{min})\|_2^2 \geq \varepsilon\Big) \\ &\leq P\Big(\frac{1}{2}\|X^T\big(X_M\beta_M^\star - X\mathrm{E}_y(\widehat{\beta}_{M(-1)})\big)\|_2^2 \geq \varepsilon\Big), \end{aligned} \tag{10}$$

where $\widehat{\beta}_{M(-1)}$ is the least squares estimate corresponding to the subset of covariates M with one covariate removed so that

$$\|\mathrm{E}_y(\widehat{\beta}_{M(-1)})\|_0 \leq |M| - 1.$$

The covariate removed is chosen to correspond to the smallest (in magnitude) component of $\widehat{\beta}_M$. To refine the bound on $\mathrm{E}[h_M^\varepsilon(\beta_M^\star)]$, decompose the last expression in (10) using the triangle inequality as follows.

$$\begin{aligned} &P\Big(\frac{1}{2}\|X^T(X_M\beta_M^\star - X_{M(-1)}\mathrm{E}_y(\widehat{\beta}_{M(-1)}))\|_2^2 \geq \varepsilon\Big) \\ &\quad\leq P\Big(\frac{1}{2}\|X^T X_M(\beta_M^\star - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\Big) \\ &\qquad + 1\Big\{\frac{1}{2}\|X^T X_M\big(\widehat{\beta}_M - \mathrm{E}_y(\widehat{\beta}_M)\big)\|_2^2 \geq \frac{\varepsilon}{9}\Big\}, \\ &\qquad + 1\Big\{\frac{1}{2}\|X^T(H_M - H_{M(-1)})\mu_T\|_2^2 \geq \frac{\varepsilon}{9}\Big\}, \end{aligned} \tag{11}$$

where $H_M = X_M(X_M^T X_M)^{-1}X_M^T$, and $\mu_T = X_{M_T}\mathrm{E}(\widehat{\beta}_{M_T})$. Observe that by Lemma 3.5 the first term on the right hand side,

$$P\Big(\frac{1}{2}\|X^T X_M(\beta^\star_M - \widehat{\beta}_M)\|_2^2 \geq \frac{\varepsilon}{9}\Big) \leq \frac{2^{\frac{3}{2}}3p_M\widehat{\sigma}_M\sqrt{\Lambda_M}}{\sqrt{\pi\varepsilon}(1-\frac{1}{n-p_M})}e^{-\frac{\varepsilon}{18\Lambda_M\widehat{\sigma}^2_M}}.$$

Note that the last two terms are written as an indicator function because the uncertainty here comes from $\beta^\star_M$. However, the second term results from the uncertainty in observing y. Accordingly, by Lemma 3.6,

$$P_{\mathrm{y}}\Big(1\Big\{\frac{1}{2}\|X^T X_M\big(\widehat{\beta}_M - \mathrm{E}_{\mathrm{y}}(\widehat{\beta}_M)\big)\|_2^2 \geq \frac{\varepsilon}{9}\Big\} = 0\Big) \geq 1 - \frac{3p_M\sigma^\star_M\sqrt{\Lambda_M}}{\sqrt{\varepsilon\pi}}e^{-\frac{\varepsilon}{9\sigma^{\star 2}_M\Lambda_M}}.$$

Recall that P_{y} is used to denote the probability measure associated with the distribution of y. The third term from (11) must be zero by Condition 3.3. ■

Proof of Theorem 3.8. Recall that

$$h^\varepsilon_{M_T}(\beta^\star_{M_T}) = 1\Big\{\frac{1}{2}\|X^T(X_{M_T}\beta^\star_{M_T} - Xb'_{min})\|_2^2 \geq \varepsilon_{M_T}(n,p)\Big\},$$

where b'_{min} solves

$$\min_{b\in\mathbb{R}^p}\frac{1}{2}\|X^T(X_{M_T}\beta^\star_{M_T} - Xb)\|_2^2 \quad \text{subject to} \quad \|b\|_0 \leq p_T - 1.$$

To show the desired result, let b_{min} be the solution to

$$\min_{b\in\mathbb{R}^p}\|X^T(\mu_T - Xb)\|_2^2 \quad \text{subject to} \quad \|b\|_0 \leq p_T - 1,$$

where $\mu_T = X_{M_T}\mathrm{E}(\widehat{\beta}_{M_T})$. Then observe that

$$\begin{aligned}\|X^T(\mu_T - Xb_{min})\|_2 &\leq \|X^T(\mu_T - Xb'_{min})\|_2 \\ &\leq \|X^T X_{M_T}(\mathrm{E}(\widehat{\beta}_{M_T}) - \beta^\star_{M_T})\|_2 \\ &\quad + \|X^T(X_{M_T}\beta^\star_{M_T} - Xb'_{min})\|_2.\end{aligned}$$

Note the difference between b_{min} and b'_{min} here. The rightmost term is the quantity of interest here because it will become $\mathrm{E}[h^\varepsilon_{M_T}(\beta^\star_{M_T})]$, in the next few steps. The term on the left of the inequality corresponds to the quantity in Condition 3.1.

Adding and subtracting $\widehat{\beta}_{M_T}$ inside the first term on the right side of the second inequality and applying the triangle inequality gives

$$\begin{aligned}
1\Big\{\frac{1}{2}\|X^T(\mu_T - Xb_{min})\|_2^2 \geq 9\varepsilon\Big\} \\
&\leq P\Big(\frac{1}{2}\|X^T X_{M_T}(\beta^\star_{M_T} - \widehat{\beta}_{M_T}\|_2^2 \geq \varepsilon\Big) \\
&\quad + 1\Big\{\frac{1}{2}\|X^T X_{M_T}(\widehat{\beta}_{M_T} - \mathrm{E}(\widehat{\beta}_{M_T})\|_2^2 \geq \varepsilon\Big\} \\
&\quad + \mathrm{E}[h^\varepsilon_{M_T}(\beta^\star_{M_T})],
\end{aligned}$$

and by applying Lemma 3.5,

$$\begin{aligned}
1\Big\{\frac{1}{2}\|X^T(\mu_T - Xb_{min})\|_2^2 \geq 9\varepsilon\Big\} \\
&\leq \frac{2^{\frac{3}{2}} 3 p_{M_T}\widehat{\sigma}_{M_T}\sqrt{\Lambda_{M_T}}}{\sqrt{\pi\varepsilon}(1 - \frac{1}{n-p_{M_T}})} e^{-\frac{\varepsilon}{18\Lambda_{M_T}\widehat{\sigma}^2_{M_T}}} \\
&\quad + 1\Big\{\frac{1}{2}\|X^T X_{M_T}(\widehat{\beta}_{M_T} - \mathrm{E}(\widehat{\beta}_{M_T})\|_2^2 \geq \frac{\varepsilon}{9}\Big\} \\
&\quad + \mathrm{E}[h^\varepsilon_{M_T}(\beta^\star_{M_T})]
\end{aligned}$$

where the middle term is written as an indicator function because the uncertainty here comes from $\beta^\star_{M_T}$. This indicator is 0 by Lemma 3.6 with probability exceeding (12). Therefore, since Condition 3.1 implies that the indicator on the left side of the inequality is 1,

$$1 - \frac{2^{\frac{3}{2}} 3 p_{M_T}\widehat{\sigma}_{M_T}\sqrt{\Lambda_{M_T}}}{\sqrt{\pi\varepsilon}(1 - \frac{1}{n-p_{M_T}})} e^{-\frac{\varepsilon}{18\Lambda_{M_T}\widehat{\sigma}^2_{M_T}}} \leq \mathrm{E}[h^\varepsilon_{M_T}(\beta^\star_{M_T})]$$

with probability exceeding

$$1 - \frac{3 p_{M_T}\sigma^\star_M\sqrt{\Lambda_{M_T}}}{\sqrt{\varepsilon\pi}} e^{-\frac{\varepsilon}{9\sigma^{\star 2}_M\Lambda_{M_T}}}, \tag{12}$$

due to the uncertainty in observing y. ■

The following lemma is needed for the proof of the main result, Theorem 3.9.

LEMMA B.1. *Assume all conditions and notations of Theorem 3.9. Then the following holds.*

Case 1 *This case pertains to subsets M with $|M| \leq p_T$.*

$$P\left(\bigcap_{j=1}^{p_T}\bigcap_{\mathcal{M}_j}\left\{y : \left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} \leq e^{-2p_T\log(p)}\right\}\right) \geq 1 - V_1,$$

where $\mathcal{M}_j := \{M \neq M_T : |M| = j\}$,

$$V_1 := \max_{\substack{M\neq M_T\\ |M|\leq p_T}}\left\{\frac{12p_T e^{\frac{-\Delta_M}{72}+p_T\log(p)}}{\sqrt{2\pi\Delta_M}} + p_T e^{-\frac{\Delta_M}{12}+\frac{p_T}{2}+p_T\log(p)} + p_T e^{-\frac{\xi_{n,p_T}}{48}\frac{(n-p_T-1)\Delta_M}{\log(n)^\gamma\log(p)}+\frac{n-p_T}{2}+p_T\log(p)}\right\},$$

and $\xi_{n,j} = 1 - \frac{2\log(n)^\gamma\log(p)}{(n-j-1)/2} \to 1$.

Case 2 *This case pertains to subsets M with $p_T + 1 \leq |M| \leq n^\alpha$, for any fixed $\alpha \in (0,1)$.*

$$P\left(\bigcap_{j=p_T+1}^{n^\alpha}\bigcap_{M:|M|=j}\left\{y : \left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} \leq e^{e^2(n^\alpha+j\log(p))}\right\}\right) \geq 1 - V_2,$$

where

$$V_2 := n^\alpha e^{-b_n\left[2n^\alpha - \frac{p_T}{2b_n} + \left(2-\frac{1}{b_n}-\frac{1}{2b_n\log(p)}\right)(p_T+1)\log(p)\right]} + \frac{n^\alpha e^{-\frac{n-n^\alpha-p_T}{2}+n^\alpha\log(p)}}{\sqrt{\pi(n-n^\alpha-p_T)}},$$

and $b_n := \frac{n-n^\alpha-p_T}{n-p_T-2} \to 1$.

Note that V_1 and V_2 both vanish for large n and p by Condition 3.2. This condition also ensures that $\xi_{n,j} \in (0,1)$ which is needed in the proof of this lemma.

Proof of Lemma B.1. There are two cases to consider.

Case 1 For $M \in \mathcal{M}_j$ with $j \leq p_T$ let $\xi_{n,j} = 1 - \frac{2\log(n)^\gamma\log(p)}{(n-j-1)/2} \in (0,1)$ by Condition 3.2 for large n and p. Then

(13)
$$P\left(\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}\right)$$
$$= P\left(\vec{e}^{\star T}(I_n - H_{M_T})\vec{e}^\star/\xi_{n,j} > \Delta_M + 2\sqrt{\Delta_M}Z + \vec{e}^{\star T}(I_n - H_M)\vec{e}^\star\right)$$

since by assumption $Y = \mu_T + \vec{e}^{\star}$ with $\vec{e}^{\star} \sim \mathbb{N}_n(0, I_n)$ and $\mu_T = X_{M_T}\mathrm{E}(\widehat{\beta}_{M_T})$, and so

$$\begin{aligned} RSS_{M_T} &= \vec{e}^{\star T}(I_n - H_{M_T})\vec{e}^{\star} \quad \text{and} \\ RSS_M &= \Delta_M + 2\mu_T^T(I_n - H_M)\vec{e}^{\star} + \vec{e}^{\star T}(I_n - H_M)\vec{e}^{\star} \\ &= Z_{\Delta_M} + \vec{e}^{\star T}(I_n - H_M)\vec{e}^{\star} \end{aligned}$$

where $Z_{\Delta_M} := \Delta_M + 2\sqrt{\Delta_M}Z$, and $Z \sim N(0,1)$. Recall that $\Delta_M := \mu_T^T(I_n - H_M)\mu_T$. Continuing in (13) by subtracting $\chi^2_{n-p_T}$ from both sides of the inequality gives,

$$\begin{aligned} &P\Big(\frac{RSS_{M_T}}{RSS_M} > \xi_{n,j}\Big) \\ &\qquad = P\big(\chi^2_{n-p_T}/\xi_{n,j} - \chi^2_{n-p_T} > Z_{\Delta_M} + \chi^2_{n-j} - \chi^2_{n-p_T}\big) \\ &\qquad = P\big(\chi^2_{n-p_T}(1/\xi_{n,j} - 1) > Z_{\Delta_M} + \chi^2_{p_T} - \chi^2_j\big) \\ &\qquad \le P\big(\chi^2_{n-p_T}(1/\xi_{n,j} - 1) > Z_{\Delta_M} - \chi^2_j\big) \end{aligned} \tag{14}$$

The last inequality follows because the $\chi^2_{p_T}$ random variable is non-negative, and removing it simplifies the remaining argument. Then

$$\begin{aligned} &P\Big(\frac{RSS_{M_T}}{RSS_M} > \xi_{n,j}\Big) \\ &\le P\big(\xi_{n,j}\chi^2_j + \chi^2_{n-p_T}(1-\xi_{n,j}) - 2\xi_{n,j}\sqrt{\Delta_M}Z > \xi_{n,j}\Delta_M\big) \\ &\le P\big(|Z| > \sqrt{\Delta_M}/6\big) + P\big(\chi^2_j > \Delta_M/3\big) + P\Big(\chi^2_{n-p_T} > \frac{\xi_{n,j}\Delta_M}{3(1-\xi_{n,j})}\Big). \end{aligned}$$

For the second and third term, apply the Chernoff bound, and evaluate the moment generating function for the χ^2_j and $\chi^2_{n-p_T}$ distributions at $1/4$. Accordingly,

$$\begin{aligned} &P\Big(\frac{RSS_{M_T}}{RSS_M} > \xi_{n,j}\Big) \\ &\qquad \le 2P\big(Z < -\sqrt{\Delta_M}/6\big) + e^{-\frac{\Delta_M}{12} + \frac{j}{2}} + e^{-\frac{\xi_{n,j}\Delta_M}{12(1-\xi_{n,j})} + \frac{n-p_T}{2}}. \end{aligned}$$

Finally the remaining probability can be controlled by the bound for the CDF of a standard normal random variable for $x > 0$,

$$\Phi(-x) \le \frac{\frac{1}{\sqrt{2\pi}}e^{\frac{-x^2}{2}}}{x}.$$

Hence,

$$P\Big(\frac{RSS_{M_T}}{RSS_M} > \xi_{n,j}\Big) \leq \frac{12e^{\frac{-\Delta_M}{72}}}{\sqrt{2\pi\Delta_M}} + e^{-\frac{\Delta_M}{12}+\frac{p_T}{2}} + e^{-\frac{\xi_{n,p_T}}{48}\frac{(n-p_T-1)\Delta_M}{\log(n)^\gamma\log(p)}+\frac{n-p_T}{2}},$$

where the last inequality follows by observing that $j \leq p_T$, and recalling the expression for $\xi_{n,j}$.

Therefore, the probability that $\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}$ is satisfied for some M with $|M| \leq p_T$ is

$$P\left(\bigcup_{j=1}^{p_T}\bigcup_{\mathcal{M}_j}\left\{\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}\right\}\right) \leq \sum_{j=1}^{p_T}\binom{p}{j}\max_{\mathcal{M}_j} P\left(\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}\right).$$

Note that

$$\binom{p}{j} = \frac{p(p-1)\cdots(p-j+1)}{j!} = \frac{p^j(1-\frac{1}{p})\cdots(1-\frac{j-1}{p})}{j!} \leq p^j. \tag{15}$$

In fact, Luo and Chen (2013) show that if $\log(j)/\log(p) \to \delta$ as $p \to \infty$, for some $\delta > 0$, then $\log\binom{p}{j} = j\log(p)(1-\delta)(1+o(1))$. Thus,

$$P\left(\bigcup_{j=1}^{p_T}\bigcup_{\mathcal{M}_j}\left\{\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > \xi_{n,j}^{\frac{n-j-1}{2}}\right\}\right) \leq \sum_{j=1}^{p_T}\max_{\mathcal{M}_j}\left\{\frac{12e^{\frac{-\Delta_M}{72}+j\log(p)}}{\sqrt{2\pi\Delta_M}} + e^{-\frac{\Delta_M}{12}+\frac{p_T}{2}+j\log(p)} + e^{-\frac{\xi_{n,p_T}}{48}\frac{(n-p_T-1)\Delta_M}{\log(n)^\gamma\log(p)}+\frac{n-p_T}{2}+j\log(p)}\right\}.$$

Since $\xi_{n,p_T} \to 1$, this bound vanishes by Condition 3.2. Therefore, with probability exceeding one minus the above bound,

$$\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} \leq \left(1 - \frac{2\log(n)^\gamma\log(p)}{(n-j-1)/2}\right)^{\frac{n-j-1}{2}} \leq e^{-2\log(n)^\gamma\log(p)}$$

uniformly over all M such that $|M| \leq p_T$.

Case 2 Consider any subset M with $p_T < |M| \leq n^\alpha$ for some positive constant $\alpha < 1$, and let $\{a_n\}$ be an arbitrarily sequence of numbers. To begin, repeating the steps in (14), but subtracting $\chi^2_{n-j}/\xi_{n,j}$ on both sides instead of $\chi^2_{n-p_T}$, and replacing the label $\xi_{n,j}$ with a_n, yields

$$P\left(\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right) \leq P\big(\chi^2_j > a_n Z_{\Delta_M} + \chi^2_{n-j}(a_n - 1)\big),$$

where $Z_{\Delta_M} = \Delta_M + 2\sqrt{\Delta_M}Z$, $Z \sim N(0,1)$, and $\Delta_M = \mu_T^T(I_n - H_M)\mu_T$. Since $M_T \subset M$ implies $\Delta_M = 0$, the above bound can be simplified by including in the subset M any covariates in M_T not already included in M. Accordingly, let $M' := M \cup M_T$ which includes $j + l$ covariates, where $l \in \{0, \ldots, p_T\}$ is the number of covariates not shared by M and M_T. Then because $RSS_{M'} \leq RSS_M$,

$$\begin{aligned} P\left(\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right) &\leq P\left(\left(\frac{RSS_{M_T}}{RSS_{M'}}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right) \\ &\leq P\big(\chi^2_{j+l} > \chi^2_{n-j-l}(a_n - 1)\big), \end{aligned}$$

and for any nonnegative $s \in \mathbb{R}$,

$$\begin{aligned} &P\Big(\frac{RSS_{M_T}}{RSS_M} > a_n\Big) \\ &\leq P\big(\{\chi^2_{j+l} > s(a_n - 1)\} \cap \{\chi^2_{n-j-l} \geq s\}\big) \\ &\qquad + P\big(\{\chi^2_{j+l} > \chi^2_{n-j-l}(a_n - 1)\} \cap \{\chi^2_{n-j-l} < s\}\big) \\ &\leq P\big(\chi^2_{j+l} > s(a_n - 1)\big) + P(\chi^2_{n-j-l} < s) \end{aligned} \tag{16}$$

Consider each of these last two terms in turn. For the first term apply the Chernoff bound, and evaluate the moment generating function for the χ^2_{j+l} distribution at $1/4$. That gives

$$P\big(\chi^2_{j+l} > s(a_n - 1)\big) \leq 2^{\frac{j+l}{2}} e^{-\frac{s(a_n-1)}{4}} \leq e^{-\frac{s(a_n-1)}{4} + \frac{j+l}{2}}.$$

For the second term in (16) write out the expression to evaluate the probability explicitly, and then apply the simple bound $e^{-x} \leq 1$ for all $x \geq 0$. Noting that $s > 0$,

$$\begin{aligned} P(\chi^2_{n-j-l} < s) &\leq \frac{1}{2^{\frac{n-j-l}{2}}\Gamma\left(\frac{n-j-l}{2}\right)} \frac{s^{\frac{n-j-l}{2}}}{\frac{n-j-l}{2}} \\ &\leq \frac{(e \cdot s)^{\frac{n-j-l}{2}}}{\sqrt{2\pi}(\frac{n-j-l}{2})^{\frac{n-j-l}{2}+\frac{1}{2}}} 2^{-\frac{n-j-l}{2}}. \end{aligned}$$

where the last inequality follows from the well known Sterling lower bound on the gamma function

$$\Gamma(x) \geq \sqrt{2\pi} x^{x-\frac{1}{2}} e^{-x} \quad \text{for} \quad x > 0.$$

It is clear from the last expression that for the probability to vanish, $e \cdot s$ must grow no faster than $\frac{n-j-l}{2}$. Accordingly, choosing $s = \frac{n-j-l}{e^2}$ gives

$$P\Big(\chi^2_{n-j-l} < \frac{n-j-l}{e^2}\Big) \leq \frac{e^{-\frac{n-j-l}{2}}}{\sqrt{\pi(n-j-l)}}.$$

Combining the two bounds for (16) now yields

$$P\Big(\frac{RSS_{M_T}}{RSS_M} > a_n\Big) \leq e^{-(\frac{n-j-l}{e^2})\frac{a_n-1}{4}+\frac{j+l}{2}} + \frac{e^{-\frac{n-j-l}{2}}}{\sqrt{\pi(n-j-l)}}.$$

It only remains to choose the smallest a_n such that the first term in the bound vanishes exponentially fast so that the cumulative probability will vanish in probability over all subsets M with $|M| \leq n^\alpha$. Accordingly, it should become apparent shortly that a good choice is

$$a_n = 1 + \frac{8e^2(n^\alpha + j\log(p))}{n-j-1}. \tag{17}$$

The probability that $\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}$ is satisfied for some M with $p_T < |M| \leq n^\alpha$ is

$$P\left(\bigcup_{j=p_T+1}^{n^\alpha} \bigcup_{M:|M|=j} \left\{\Big(\frac{RSS_{M_T}}{RSS_M}\Big)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right\}\right)$$
$$\leq \sum_{j=p_T+1}^{n^\alpha} \binom{p}{j} \max_{M:|M|=j} P\left(\Big(\frac{RSS_{M_T}}{RSS_M}\Big)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right).$$

Thus, bounding the binomial coefficient as in (15), and substituting

(17) for a_n yields

$$P\left(\bigcup_{j=p_T+1}^{n^\alpha}\bigcup_{M:|M|=j}\left\{\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} > a_n^{\frac{n-j-1}{2}}\right\}\right)$$
$$\leq \sum_{j=p_T+1}^{n^\alpha} e^{-b_n\left[2n^\alpha - \frac{l}{2b_n} + \left(2-\frac{1}{b_n}-\frac{1}{2b_n\log(p)}\right)j\log(p)\right]} + \frac{e^{-\frac{n-j-p_T}{2}+j\log(p)}}{\sqrt{\pi(n-j-p_T)}}.$$
$$\leq n^\alpha e^{-b_n\left[2n^\alpha - \frac{p_T}{2b_n} + \left(2-\frac{1}{b_n}-\frac{1}{2b_n\log(p)}\right)(p_T+1)\log(p)\right]}$$
$$+ \frac{n^\alpha e^{-\frac{n-n^\alpha-p_T}{2}+n^\alpha\log(p)}}{\sqrt{\pi(n-n^\alpha-p_T)}},$$

where $b_n := \frac{n-n^\alpha-p_T}{n-p_T-2} \to 1$. Note that this bound vanishes by Condition 3.2. Therefore, with probability exceeding one minus the above bound,

$$\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} \leq \left(1+\frac{4e^2(n^\alpha+j\log(p))}{\frac{n-j-1}{2}}\right)^{\frac{n-j-1}{2}}$$
$$\leq e^{4e^2(n^\alpha+j\log(p))},$$

uniformly over all M such that $p_T < |M| \leq n^\alpha$.

■

Proof of Theorem 3.9. For $j \in \{1,\ldots,p\}$ define the following classes of subsets $\mathcal{M}_j := \{M \neq M_T : |M| = j\}$. Recall that $\widehat{\sigma}_M^2 := RSS_M/(n-p_M)$. It will first be shown that for any subset of covariates $M \neq M_T$ the ratio $\frac{r(M|y)}{r(M_T|y)}$ vanishes in probability for large n and p. Accordingly, for any $M \in \mathcal{M}_j$,

$$\frac{r(M|y)}{r(M_T|y)} = \pi^{\frac{j-p_T}{2}}\frac{\Gamma\left(\frac{n-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)}\frac{RSS_{M_T}^{\frac{n-p_T-1}{2}}}{RSS_M^{\frac{n-j-1}{2}}}\frac{\mathrm{E}[h_M^\varepsilon(\beta_M^\star)]}{\mathrm{E}[h_{M_T}^\varepsilon(\beta_{M_T}^\star)]}$$
$$= \pi^{\frac{j-p_T}{2}}\frac{\Gamma\left(\frac{n-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)}\left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}}RSS_{M_T}^{\frac{j-p_T}{2}}\frac{\mathrm{E}[h_M^\varepsilon(\beta_M^\star)]}{\mathrm{E}[h_T^\varepsilon(\beta_{M_T}^\star)]}$$

Before proceeding with the rest of the proof, a the following notation is needed. V_1 and V_2 are as stated in Lemma B.1. As in Theorem 3.8, define

$$V_3 := \frac{3p_{M_T}\sigma_M^\star\sqrt{\Lambda_{M_T}}}{\sqrt{\varepsilon_{M_T}(n,p)\pi}}e^{-\frac{\varepsilon_{M_T}(n,p)}{9\sigma_M^{\star 2}\Lambda_{M_T}}},$$

and corresponding to Theorem 3.7, define

$$V_4 := \sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{3p_M \sigma_M^\star \sqrt{\Lambda_M}}{\sqrt{\varepsilon_M(n,p)\pi}} e^{-\frac{\varepsilon_M(n,p)}{9\sigma_M^{\star 2}\Lambda_M}} \leq \sum_{j=1}^{n^\alpha} \max_{\mathcal{M}_j} \frac{3p_M \sigma_M^\star \sqrt{\Lambda_M}}{\sqrt{\varepsilon_M(n,p)\pi}} e^{-\frac{\varepsilon_M(n,p)}{9\sigma_M^{\star 2}\Lambda_M} + j\log(p)}$$

by bounding the binomial coefficient as in (15). Note that V_4 then vanishes by Condition 3.4. Further, recall that $RSS_{M_T} \sim \chi^2_{n-p_T}$, so by the Chernoff bound (evaluating the moment generating function at $1/4$),

$$P\big(\chi^2_{n-p_T} > 3(n-p_T)\big) \leq \underbrace{e^{-\frac{n-p_T}{4}}}_{=: V_5}.$$

With this notation it is now possible to account for all of the uncertainty due to y.

Accordingly, by Theorem 3.8, with probability exceeding $1 - V_3 - V_5$,

$$\frac{r(M|y)}{r(M_T|y)} \leq \frac{\Gamma\left(\frac{n-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)} \left(\frac{RSS_{M_T}}{RSS_M}\right)^{\frac{n-j-1}{2}} \frac{(3\pi(n-p_T))^{\frac{j-p_T}{2}} \mathrm{E}[h^\varepsilon_M(\beta^\star_M)]}{1 - g_1(M_T,n,p)}, \tag{18}$$

where

$$g_1(M,n,p) := \frac{2^{\frac{3}{2}} 3j\widehat{\sigma}_M \sqrt{\Lambda_M}}{\sqrt{\pi\varepsilon_M(n,p)}(1-\frac{1}{n-j})} e^{-\frac{\varepsilon_M(n,p)}{18\Lambda_M \widehat{\sigma}^2_M}}.$$

Further, fix $A_1 \in (0,1)$, and by Condition 3.4 choose n and p sufficiently large so that $g_1(M_T,n,p) < A_1$.

Suppose $j \leq p_T$. As in Jameson (2013), the ratio of gamma functions can be bounded by

$$\frac{\Gamma\left(\frac{n-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)} = \frac{\Gamma\left(\frac{n-p_T}{2} + \frac{p_T-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)} \leq \left(\frac{n-p_T}{2}\right)\left(\frac{n-j}{2}\right)^{\frac{p_T-j}{2}-1}.$$

Applying Lemma B.1 to bound the ratio of RSS, and bounding the expectation by 1, with probability exceeding $1 - V_1 - V_3 - V_5$, (18) implies

$$\begin{aligned} \frac{r(M|y)}{r(M_T|y)} &\leq \left(\frac{n-p_T}{2}\right)\left(\frac{n-j}{2}\right)^{\frac{p_T-j}{2}-1} \frac{(3\pi(n-p_T))^{\frac{j-p_T}{2}}}{(1-A_1)e^{2\log(n)^\gamma \log(p)}} \\ &\leq \frac{e^{-2\log(n)^\gamma \log(p)}}{1-A_1}, \end{aligned} \tag{19}$$

where the last inequality follows for all $n \geq \frac{p_T}{1-\frac{1}{6\pi}}$. Since $n >> p_T$, assume without loss of generality that n is sufficiently large.

Conversely, suppose $p_T < j \leq n^\alpha$ for any fixed positive constant $\alpha < 1$. In this setting, as in Jameson (2013), the ratio of gamma functions can be bounded by

$$\left(\frac{\Gamma\left(\frac{n-j}{2}\right)}{\Gamma\left(\frac{n-p_T}{2}\right)}\right)^{-1} = \frac{\Gamma\left(\frac{n-p_T}{2}\right)}{\Gamma\left(\frac{n-j}{2}\right)} = \frac{\Gamma\left(\frac{n-j}{2}+\frac{j-p_T}{2}\right)}{\Gamma\left(\frac{n-j}{2}\right)} \geq \left(\frac{n-j}{2}-1\right)^{\frac{j-p_T}{2}}.$$

Applying Theorem 3.7 to bound the expectation, and applying Lemma B.1 to bound the ratio of RSS, with probability exceeding $1 - V_2 - V_3 - V_4 - V_5$, (18) implies

$$\begin{aligned}\frac{r(M|y)}{r(M_T|y)} &\leq \frac{e^{4e^2(n^\alpha + j\log(p))}}{\left(\frac{n-j}{2}-1\right)^{\frac{j-p_T}{2}}}(3\pi(n-p_T))^{\frac{j-p_T}{2}}\frac{g_1(M,n,p)}{1-A_1} \\ &\leq \frac{2^{\frac{3}{2}}3j\widehat{\sigma}_M\sqrt{\Lambda_M}e^{-\frac{\varepsilon_M(n,p)}{18\Lambda_M\widehat{\sigma}_M^2}+4e^2(n^\alpha+j\log(p))+\frac{j-p_T}{2}\log\left(\frac{6\pi}{1-\frac{n^\alpha+2}{n}}\right)}}{\sqrt{\pi\varepsilon_M(n,p)}(1-\frac{1}{n-j})(1-A_1)}\end{aligned} \tag{20}$$

Notice that Theorem 3.7 is being applied here with n^α in place of $c_1 n$. This can be done without loss of generality because n^α grows slower than $c_1 n$, for any choices of $\alpha, c_1 < 1$.

It will now be shown that $\sum_{j=1}^{n^\alpha}\sum_{M\in\mathcal{M}_j}\frac{r(M|y)}{r(M_T|y)}$ vanishes in probability for large n and p. Apply the bounds in (19) and (20) in the following argument.

$$\sum_{j=1}^{n^\alpha}\sum_{M\in\mathcal{M}_j}\frac{r(M|y)}{r(M_T|y)} \leq \underbrace{\sum_{j=1}^{p_T}\binom{p}{j}\max_{M\in\mathcal{M}_j}\frac{r(M|y)}{r(M_T|y)}}_{=:S_1} + \underbrace{\sum_{j=p_T+1}^{n^\alpha}\binom{p}{j}\max_{|M|=j}\frac{r(M|y)}{r(M_T|y)}}_{=:S_2}$$

Consider S_1 and S_2 separately.

With probability exceeding $1 - V_1 - V_3 - V_5$,

$$S_1 \leq \frac{1}{1-A_1}\sum_{j=1}^{p_T}e^{-2\log(n)^\gamma\log(p)+j\log(p)} \leq \frac{p_T e^{-(2\log(n)^\gamma - p_T)\log(p)}}{1-A_1},$$

by bounding the binomial coefficient as in (15), and with probability exceed-

ing $1 - V_2 - V_3 - V_4 - V_5$,

$$S_2 \leq A_2 \sum_{j=p_T+1}^{n^\alpha} \frac{\frac{j}{1-\frac{1}{n-j}} \max_{|M|=j} \left(\frac{\widehat{\sigma}_M^2 \Lambda_M}{\varepsilon_M(n,p)} \right)^{\frac{1}{2}}}{e^{\min_{|M|=j} \left\{ \frac{\varepsilon_M(n,p)}{18\Lambda_M \widehat{\sigma}_M^2} - 4e^2(n^\alpha + j\log(p)) - \frac{j-p_T}{2} \log\left(\frac{6\pi}{1-\frac{n^\alpha+2}{n}} \right) - j\log(p) \right\}}}$$
$$\leq \frac{A_2}{n^{1-\alpha}} \cdot \max_{p_T < |M| \leq n^\alpha} \left(\frac{\widehat{\sigma}_M^2 \Lambda_M}{\varepsilon_M(n,p)} \right)^{\frac{1}{2}},$$

for some positive constant A_2. The last inequality follows by Condition 3.4.

Thus, for sufficiently large n and p, with probability exceeding $1 - V_1 - V_2 - V_3 - V_4 - V_5$,

$$\sum_{j=1}^{n^\alpha} \sum_{M \in \mathcal{M}_j} \frac{r(M|y)}{r(M_T|y)} \leq \frac{p_T e^{-(2\log(n)^\gamma - p_T)\log(p)}}{1 - A_1} + \frac{A_2}{n^{1-\alpha}} \max_{p_T < |M| \leq n^\alpha} \left(\frac{\widehat{\sigma}_M^2 \Lambda_M}{\varepsilon_M(n,p)} \right)^{\frac{1}{2}}$$

which by Conditions 3.1-3.4 vanishes for large n and p. The proof is now complete by noticing that

$$r(M_T|y) = \frac{r(M_T|y)}{\sum_{j=1}^{n^\alpha} \sum_{M:|M|=j} r(M|y)} = \frac{1}{1 + \sum_{j=1}^{n^\alpha} \sum_{\mathcal{M}_j} \frac{r(M|y)}{r(M_T|y)}}.$$

■

E-MAIL: jpwill@live.unc.edu